arXiv:1608.08704v2 [cs.LO] 1 Sep 2016

Near-Optimal Lower Bounds on Quantifier Depth and Weisfeiler–Leman Refinement Steps*

Christoph Berkholz
Humboldt-Universität zu Berlin

Jakob Nordström
KTH Royal Institute of Technology

September 2, 2016

Abstract

We prove near-optimal trade-offs for quantifier depth versus number of variables in first-order logic by exhibiting pairs of n-element structures that can be distinguished by a k-variable first-order sentence but where every such sentence requires quantifier depth at least $n^{\Omega(k/\log k)}$. Our trade-offs also apply to first-order counting logic, and by the known connection to the k-dimensional Weisfeiler–Leman algorithm imply near-optimal lower bounds on the number of refinement iterations.

A key component in our proof is the hardness condensation technique recently introduced by [Razborov '16] in the context of proof complexity. We apply this method to reduce the domain size of relational structures while maintaining the minimal quantifier depth to distinguish them in finite variable logics.

1 Introduction

The k-variable fragment of first-order logic L^k consists of those first-order sentences that use at most k different variables. A simple example is the L^2 sentence

$$\exists x \exists y (Exy \wedge \exists x (Eyx \wedge \exists y (Exy \wedge \exists x Eyx))) \tag{1.1}$$

stating that there exists a directed path of length 4 in a digraph. Extending L^k with counting quantifiers $\exists^{\geq i} x$ yields C^k, which can be more economical in terms of variables. As an illustration, the L^8 sentence

$$\exists x \exists y_1 \cdots \exists y_7 \big(\textstyle\bigwedge_{i \neq j} y_i \neq y_j \wedge \bigwedge_i Exy_i \big) \tag{1.2}$$

stating the existence of a vertex of degree at least 7 in a graph can be written more succinctly as the C^2 sentence

$$\exists x \exists^{\geq 7} y Exy \ . \tag{1.3}$$

Bounded variable fragments of first order logic have found numerous applications in finite model theory and related areas (see [Gro98] for a survey). Their importance stems from the fact that the model checking problem (given a finite relational structure $\mathcal{A}$ and a sentence φ, does $\mathcal{A}$ satisfy φ?) can be decided in polynomial time [Imm82, Var95]. Moreover, the equivalence problem (given two finite relational structures $\mathcal{A}$ and $\mathcal{B}$, do they satisfy the same sentences?) for L^k and C^k can be decided in time $n^{O(k)}$ [IL90], i.e., polynomial for constant k.

*This is the full-length version of a paper with the same title which appeared in *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '16)*.

1.1 Quantifier Depth

If $\mathcal{A}$ and $\mathcal{B}$ are not equivalent in L^k or C^k, then there exists a sentence φ that defines a distinguishing property, i.e., such that $\mathcal{A} \models \varphi$ and $\mathcal{B} \not\models \varphi$, which certifies that the structures are non-isomorphic. But how complex can such a sentence be? In particular, what is the minimal quantifier depth of an L^k or C^k sentence that distinguishes two n-element relational structures $\mathcal{A}$ and $\mathcal{B}$? The best upper bound for the quantifier depth of L^k and C^k is n^{k-1} [IL90], while to the best of our knowledge the strongest lower bounds have been only linear in n [CFI92, Gro99, Fü01]. In this paper we present a near-optimal lower bound of $n^{\Omega(k/\log k)}$.

Theorem 1.1. *There exist $\varepsilon > 0$, $K_0 \in \mathbb{N}$ such that for all k, n with $K_0 \leq k \leq n^{1/12}$ there is a pair of n-element $(k-1)$-ary relational structures $\mathcal{A}_n, \mathcal{B}_n$ that can be distinguished in k-variable first-order logic but satisfy the same L^k and C^k sentences up to quantifier depth $n^{\varepsilon k/\log k}$.*

Note that any two non-isomorphic n-element σ-structures $\mathcal{A}$ and $\mathcal{B}$ can always be distinguished by a simple n-variable first-order sentence of quantifier depth n, namely

$$\exists x_1 \cdots \exists x_n \left(\bigwedge_{i \neq j} x_i \neq x_j \;\wedge \bigwedge_{\substack{R \in \sigma, \\ (v_{i_1}, \ldots, v_{i_r}) \in R^{\mathcal{A}}}} R x_{i_1}, \ldots, x_{i_r} \;\wedge \bigwedge_{\substack{R \in \sigma, \\ (v_{i_1}, \ldots, v_{i_r}) \notin R^{\mathcal{A}}}} \neg R x_{i_1}, \ldots, x_{i_r} \right) . \tag{1.4}$$

Since our $n^{\Omega(k/\log k)}$ lower bound for k-variable logics grows significantly faster than this trivial upper bound n on the quantifier depth as the number of variables increases, Theorem 1.1 also describes a trade-off in the super-critical regime above worst-case investigated by Razborov [Raz16]: If one reduces one complexity measure (the number of variables), then the other complexity parameter (the quantifier depth) increases sharply even beyond its worst-case upper bound.

The equivalence problem for C^{k+1} is known to be closely related to the *k-dimensional Weisfeiler–Leman algorithm* (k-WL) for testing non-isomorphism of graphs and, more generally, relational structures. It was shown by Cai, Fürer, and Immerman [CFI92] that two structures are distinguished by k-WL if and only if there exists a C^{k+1} sentence that differentiates between them. Moreover, the quantifier depth of such a sentence also relates to the complexity of the WL algorithm in that the number of iterations k-WL needs to tell $\mathcal{A}$ and $\mathcal{B}$ apart coincides with the minimal quantifier depth of a distinguishing C^{k+1} sentence. Therefore, Theorem 1.1 also implies a near-optimal lower bound on the number of refinement steps required in the Weisfeiler–Leman algorithm. We discuss this next.

1.2 The Weisfeiler–Leman Algorithm

The Weisfeiler–Leman algorithm, independently introduced by Babai in 1979 and by Immerman and Lander in [IL90] (cf. [CFI92] and [Bab16] for historic notes), is a hierarchy of methods for isomorphism testing that iteratively refine a partition (or colouring) of the vertex set, ending with a *stable colouring* that classifies *similar vertices*. Since no isomorphism can map non-similar vertices to each other, this reduces the search space. Moreover, if two structures end up with different stable colourings, then we can immediately deduce that the structures are non-isomorphic. The 1-dimensional Weisfeiler–Leman algorithm, better known as *colour refinement*, initially colours the vertices according to their degree (clearly, no isomorphism identifies vertices of different degree). The vertex colouring is then refined based on the colour classes of the neighbours. For example, two degree-5 vertices get different colours in the next step if they have a different number of degree-7 neighbours. This refinement step is repeated until the colouring stays stable (i.e., every pair of equally coloured vertices have the same number of neighbours in every other colour class). This algorithm is already quite strong and is extensively used in practical graph isomorphism algorithms.

In k-dimensional WL this idea is generalized to colourings of *k-tuples* of vertices. Initially the k-tuples are coloured by their isomorphism type, i.e., two tuples $\vec{v} = (v_1, \ldots, v_k)$ and $\vec{w} = (w_1, \ldots, w_k)$ get different colours if the mapping $v_i \mapsto w_i$ is not an isomorphism on the substructures induced on

$\{v_1, \dots, v_k\}$ and $\{w_1, \dots, w_k\}$. In the refinement step, we consider for each k-tuple $\vec{v} = (v_1, \dots, v_k)$ and every vertex v the colours of the tuples $\vec{v}_j := (v_1, \dots, v_{j-1}, v, v_{j+1}, \dots, v_k)$, where v is substituted at the jth position in the tuple $\vec{v}$. We refer to the tuple $(c(\vec{v}_1), \dots, c(\vec{v}_k))$ of these k colours as the *colour type* $\mathtt{t}(\vec{v}, v)$ and let v be a $\mathtt{t}$-neighbour of $\vec{v}$ if $\mathtt{t} = \mathtt{t}(\vec{v}, v)$. Now two tuples $\vec{v}$ and $\vec{w}$ get different colours if they are already coloured differently, or if there exists a colour type $\mathtt{t}$ such that $\vec{v}$ and $\vec{w}$ have a different number of $\mathtt{t}$-neighbours. The refinement step is repeated until the colouring stays stable. Since in every round the number of colour classes grows, the process stops after at most n^k steps. The colour names can be chosen in such a way that the stable colouring is canonical, which means that two isomorphic structures end up with the same colouring, and such a canonical stable colouring can be computed in time $n^{O(k)}$.

This simple combinatorial algorithm is surprisingly powerful. Grohe [Gro12] showed that for every nontrivial graph class that excludes some minor (such as planar graphs or graphs of bounded treewidth) there exists some k such that k-WL computes a different colouring for all non-isomorphic graphs, and hence solves graph isomorphism in polynomial time on that graph class. Weisfeiler–Leman has also been used as a subroutine in algorithms that solve graph isomorphism on all graphs. As one part of his very recent graph isomorphism algorithm, Babai [Bab16] applies k-WL for polylogarithmic k to relational (k-ary) structures and makes use of the quasi-polynomial running time of this algorithm.

Given the importance of the Weisfeiler–Leman procedure, it is a natural question to ask whether the trivial n^k upper bound on the number of refinement steps is tight. By the correspondence between the number of refinement steps of k-WL and the quantifier depth of C^{k+1} [CFI92], our main result implies a near-optimal lower bound even up to polynomial, but still sublinear, values of k (i.e., $k = n^\delta$ for small enough constant δ).

Theorem 1.2. *There exist $\varepsilon > 0$, $K_0 \in \mathbb{N}$ such that for all k, n with $K_0 \leq k \leq n^{1/12}$ there is an n-element k-ary relational structure $\mathcal{A}_n$ for which the k-dimensional Weisfeiler–Leman algorithm needs $n^{\varepsilon k/\log k}$ refinement steps to compute the stable colouring.*

In addition to the near-optimal lower bounds for a specific dimension (or number of variables) k, we also obtain the following trade-off between the dimension and the number of refinement steps: If we fix two parameters ℓ_1 and ℓ_2 (possibly depending on n) satisfying $\ell_1 \leq \ell_2 \leq n^{1/6}/\ell_1$, then there are n-element structures such that k-WL needs $n^{\Omega(\ell_1/\log \ell_2)}$ refinement steps for all $\ell_1 \leq k \leq \ell_2$. A particularly interesting choice of parameters is $\ell_1 = \log^c n$ for some constant $c > 1$ and $\ell_2 = n^{1/7}$. This implies the following quasi-polynomial lower bound on the number of refinement steps for Weisfeiler–Leman from polylogarithmic dimension all the way up to dimension $n^{1/7}$.

Theorem 1.3. *For every $c > 1$ there is a sequence of n-element relational structures $\mathcal{A}_n$ for which the k-dimensional Weisfeiler–Leman algorithm needs $n^{\Omega(\log^{c-1} n)}$ refinement steps to compute the stable colouring for all k with $\log^c n \leq k \leq n^{1/7}$.*

1.3 Previous Lower Bounds

In their seminal work [CFI92], Cai, Fürer and Immerman established the existence of non-isomorphic n-vertex graphs that cannot be distinguished by any first-order counting sentence with $o(n)$ variables. Since every pair of non-isomorphic n-element structures can be distinguished by a C^n (or even L^n) sentence (as shown in (1.4) above), this result also implies a linear lower bound on the quantifier depth of C^k if $k = \Omega(n)$. For all constant $k \geq 2$, a linear $\Omega(n)$ lower bound on the quantifier depth of C^k follows implicitly from an intricate construction of Grohe [Gro99], which was used to show that the equivalence problems for L^k and C^k are complete for polynomial time. An explicit linear lower bound based on a simplified construction was subsequently presented by Fürer [Fü01].

For the special case of C^2, Krebs and Verbitsky [KV15] recently obtained an improved $(1 - o(1))n$ lower bound on the quantifier depth, nearly matching the upper bound n. In contrast, Kiefer and Schweitzer [KS16] showed that if two n-vertex graphs can be distinguished by a C^3 sentence, then

there is always a distinguishing sentence of quantifier depth $O(n^2/\log n)$. Hence, the trivial n^2 upper bound is not tight in this case.

As far as we are aware, the current paper presents the first lower bounds that are super-linear in the domain size n.

1.4 Discussion of Techniques

The hard instances we construct are based on propositional XOR (exclusive or) formulas, which can alternatively be viewed as systems of linear equations over $\mathrm{GF}(2)$. There is a long history of using XOR formulas for proving lower bounds in different areas of theoretical computer science such as, e.g., finite model theory, proof complexity, and combinatorial optimization/hardness of approximation. Our main technical insight is to combine two methods that, to the best of our knowledge, have not been used together before, namely Ehrenfeucht-Fraïssé games on structures based on XOR formulas and hardness amplification by variable substitution.

More than three decades ago, Immerman [Imm81] presented a way to encode an XOR formula into two graphs that are isomorphic if and only if the formula is satisfiable. This can then be used to show that the two graphs cannot be distinguished by a sentence with few variables or low quantifier depth using Ehrenfeucht-Fraïssé games. Arguably the most important application of this method is the result in [CFI92] establishing that a linear number of variables is needed to distinguish two graphs in first-order counting logic. Graph constructions based on XOR formulas have also been used to prove lower bounds on the quantifier depth of C^k [Imm81, Fü01]. We remark that for our result we have to use a slightly different encoding of XOR formulas into relational structures rather than graphs.

In proof complexity, various flavours of XOR formulas (usually called *Tseitin formulas* when used to encode the *handshaking lemma* saying that the sum of all vertex degrees in an undirected graph has to be an even number) have been employed to obtain lower bounds for proof systems such as resolution [Urq87], polynomial calculus [BGIP01], and bounded-depth Frege [Ben02a]. Such formulas have also played an important role in many lower bounds for the Positivstellensatz/sums-of-squares proof system [Gri01, KI06, Sch08] corresponding to the Lasserre semidefinite programming hierarchy, which has been the focus of much recent interest in the context of combinatorial optimization.[1] Another use of XOR in proof complexity has been for hardness amplification, where one takes a (typically non-XOR) formula that is moderately hard with respect to some complexity measure, substitutes all variables by exclusive ors over pairwise distinct sets of variables, and then shows that the new *XORified* formula must be very hard with respect to some other (more important) complexity measure. This technique was perhaps first made explicit in [Ben02b] (attributed there to personal communication with Michael Alekhnovich and Alexander Razborov, with a note that it is also very similar in spirit to an approach used in [BW01]) and has later appeared in, e.g., [BP07, BN08, BN11, BNT13, FLM$^+$13].

An even more crucial role in proof complexity is played by well-connected so-called *expander graphs*. For instance, given a formula in conjunctive normal form (CNF) one can look at its bipartite clause-variable incidence graph (CVIG), or some variant of the CVIG derived from the combinatorial structure of the formula, and prove that if this graph is an expander, then this implies that the formula must be hard for proof systems such as resolution [BW01] and polynomial calculus [AR03, MN15].

In a striking recent paper [Raz16], Razborov combines XORification and expansion in a simple (with hindsight) but amazingly powerful way. Namely, instead of replacing every variable by an XOR over new, fresh variables, he recycles variables from a much smaller pool, thus decreasing the total number of variables. This means that the hardness amplification proofs no longer work, since they crucially use that all new substitution variables are distinct. But here expansion come into play. If the pattern of variable substitutions is described by a strong enough bipartite expander, it turns out that locally there is enough "freshness" even among the recycled variables to make the hardness amplification go through over a fairly wide range of the parameter space. And since the formula has not only become harder but has also

[1] No proof complexity is needed in this paper, and so readers unfamiliar with these proof systems need not worry—this is just an informal overview.

had the number of variables decreased, this can be viewed as a kind of *hardness compression* or *hardness condensation*.

What we do in this paper is to first revisit Immerman's old quantifier depth lower bound for first-order counting logic [Imm81] and observe that the construction can be used to obtain an improved scalable lower bound for the k-variable fragment. We then translate Razborov's hardness condensation technique [Raz16] into the language of finite variable logics and use it—perhaps somewhat amusingly applied to XORification of XOR formulas, which is usually not the case in proof complexity—to reduce the domain size of relational structures while maintaining the minimal quantifier depth required to distinguish them.

1.5 Outline of This Paper

The rest of this paper is organized as follows. In Section 2 we describe how to translate XOR formulas to relational structures and play combinatorial games on these structures. This then allows us to state our main technical lemmas in Section 3 and show how these lemmas yield our results. Turning to the proofs of these technical lemmas, in Section 4 we present a version of Immerman's quantifier depth lower bound for XOR formulas, and in Section 5 we apply Razborov's hardness condensation technique to these formulas. Finally, in Section 6 we make some concluding remarks and discuss possible directions for future research. Some proofs of technical results needed in the paper are deferred to Appendix A.

2 From XOR Formulas to Relational Structures

In this paper all structures are finite and defined over a relational signature σ. We use the letters X, E, and R for unary, binary, and r-ary relation symbols, respectively, and let $X^{\mathcal{A}}$, $E^{\mathcal{A}}$, and $R^{\mathcal{A}}$ be their interpretation in a structure $\mathcal{A}$. We write $V(\mathcal{A})$ to denote the domain of the structure $\mathcal{A}$. The *k-variable fragment of first-order logic* L^k consists of all first-order formulas that use at most k different variables (possibly re-quantifying them as in Equation (1.1)). We also consider *k-variable first-order counting logic* C^k, which is the extension of L^k by counting quantifiers $\exists^{\geq i} x \varphi(x)$, stating that there exist at least i elements $u \in V(\mathcal{A})$ such that $(\mathcal{A}, u) \models \varphi(x)$. For a survey of finite variable logics and their applications we refer the reader to, e.g., [Gro98].

An *ℓ-XOR clause* is a tuple $(x_1, \dots, x_\ell, a)$ consisting of ℓ Boolean variables and a Boolean value $a \in \{0,1\}$. We refer to ℓ as the *width* of the clause. An assignment α *satisfies* $(x_1, \dots, x_\ell, a)$ if $\alpha(x_1) + \cdots + \alpha(x_\ell) \equiv a \pmod 2$. An ℓ-XOR formula F is a conjunction of XOR clauses of width at most ℓ and is satisfied by an assignment α if α satisfies all clauses in F.

For every ℓ-XOR formula F on n variables we can define a pair of $2n$-element structures $\mathcal{A} = \mathcal{A}(F)$ and $\mathcal{B} = \mathcal{B}(F)$ that are isomorphic if and only if F is satisfiable. The domain of the structures contains two elements x_i^0 and x_i^1 for each Boolean variable x_i. There is one unary predicate X_i for every variable x_i identifying the corresponding two elements x_i^0 and x_i^1. Hence these unary relations partition the domain of the structures into two-element sets, i.e., $X_i^{\mathcal{A}} = X_i^{\mathcal{B}} = \{x_i^0, x_i^1\}$. To encode the XOR clauses, we introduce one m-ary relation R_m for every $1 \leq m \leq \ell$ and set

$$R_m^{\mathcal{A}} = \left\{ \left(x_{i_1}^{a_1}, \dots, x_{i_m}^{a_m}\right) \middle| (x_{i_1}, \dots, x_{i_m}, a) \in F,\ \textstyle\sum_i a_i \equiv 0 \pmod 2 \right\} \tag{2.1a}$$

and

$$R_m^{\mathcal{B}} = \left\{ \left(x_{i_1}^{a_1}, \dots, x_{i_m}^{a_m}\right) \middle| (x_{i_1}, \dots, x_{i_m}, a) \in F,\ \textstyle\sum_i a_i \equiv a \pmod 2 \right\} . \tag{2.1b}$$

Every bijection β between the domains of $\mathcal{A}(F)$ and $\mathcal{B}(F)$ that preserves the unary relations X_i can be translated to an assignment α for the XOR formula via the correspondence

$$\alpha(x_i) = 0 \Leftrightarrow \beta(x_i^0) = x_i^0 \Leftrightarrow \beta(x_i^1) = x_i^1 \tag{2.2a}$$

and

$$\alpha(x_i) = 1 \Leftrightarrow \beta(x_i^0) = x_i^1 \Leftrightarrow \beta(x_i^1) = x_i^0 \ . \tag{2.2b}$$

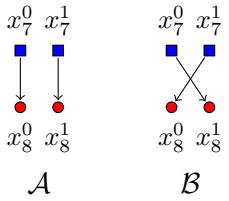

Figure 1: Structure encoding of $F = \{(x_7, x_8, 1)\}$.

Moreover, it is not hard to show that such a bijection defines an isomorphism between $\mathcal{A}(F)$ and $\mathcal{B}(F)$ if and only if the corresponding assignment satisfies F. See Figure 1 for a small example illustrating the construction.

This kind of encodings of XOR formulas into relational structures have been very useful for proving lower bounds for finite variable logics in the past. Our transformation of XOR clauses of width ℓ into ℓ-ary relational structures resembles the way Gurevich and Shelah [GS96] encode XOR formulas as hypergraphs. It is also closely related to the way Cai, Fürer, and Immerman [CFI92] obtain two non-isomorphic graphs $\mathcal{G}$ and $\mathcal{H}$ from an unsatisfiable 3-XOR formula F in the sense that $\mathcal{G}$ and $\mathcal{H}$ can be seen to be the incidence graphs of our structures $\mathcal{A}(F)$ and $\mathcal{B}(F)$.

In order to prove our main result, we make use of the combinatorial characterization of quantifier depth of finite-variable logics in terms of pebble games for L^k and C^k, which are played on two given relational structures. Since in our case the structures are based on XOR formulas, for convenience we will consider a simplified combinatorial game that is played directly on the XOR formulas rather than on their structure encodings. We first describe this game and then show in Lemma 2.1 that this yields an equivalent characterization.

The *r-round k-pebble game* is played on an XOR formula F by two players, whom we will refer to as Player 1 and Player 2. A position in the game is a partial assignment α of at most k variables of F and the game starts with the empty assignment. In each round, Player 1 can delete some variable assignments from the current position (he chooses some $\alpha' \subseteq \alpha$). If the current position assigns values to exactly k variables, then Player 1 has to delete at least one variable assignment. Afterwards, Player 1 chooses some currently unassigned variable x and asks for its value. Player 2 answers by either 0 or 1 (independently of any previous answers to the same question) and adds this variable assignment to the current position.

A winning position for Player 1 is an assignment falsifying some clause from F. Player 1 wins the r-round k-pebble game if he has a strategy to win every play of the k-pebble game within at most r rounds. Otherwise, we say that Player 2 wins (or survives) the r-round k-pebble game. Player 1 *wins the k-pebble game* if he wins the r-round k-pebble game within a finite number of rounds r. Note that if Player 1 wins the k-pebble game, then he can always win the k-pebble within $2^k n^{k+1}$ rounds, because there are at most $\sum_{i=0}^{k} 2^i \binom{n}{i} \leq 2^k n^{k+1}$ different positions with at most k pebbles on n-variable XOR formulas. We say that Player 1 *can reach a position* β from a position α within r rounds if he has a strategy such that in every play of the r-round k-pebble game starting from position α he either wins or ends up with position β.

As a side remark, we note that if we expand the XOR formula to CNF, then our pebble game is the same as the so-called *Boolean existential pebble game* played on this CNF encoding and therefore also characterizes the resolution width required for the corresponding CNF formula as shown in [AD08]. Intuitively, it is this correspondence that enables us to apply the proof complexity techniques from [Raz16] in our setting. We will not need to use any concepts from proof complexity in this paper, however, but will present a self-contained proof, and so we do not elaborate further on this connection.

Let us now show that the game described above is equivalent to the pebble game for L^k and to the bijective pebble game for C^k played on the structures $\mathcal{A}(F)$ and $\mathcal{B}(F)$.

Lemma 2.1. *Let k, p, r be integers such that $r > 0$ and $k \geq p$ and let F be a p-XOR formula giving rise to structures $\mathcal{A} = \mathcal{A}(F)$ and $\mathcal{B} = \mathcal{B}(F)$ as described in the paragraph preceding* (2.1a)–(2.1b)*. Then*

the following statements are equivalent:

(a) Player 1 wins the r-round k-pebble game on F.

(b) There is a k-variable first-order sentence $\varphi \in \mathsf{L}^k$ of quantifier depth r such that $\mathcal{A}(F) \models \varphi$ and $\mathcal{B}(F) \not\models \varphi$.

(c) There is a k-variable sentence in first-order counting logic $\varphi \in \mathsf{C}^k$ of quantifier depth r such that $\mathcal{A}(F) \models \varphi$ and $\mathcal{B}(F) \not\models \varphi$.

(d) The $(k-1)$-dimensional Weisfeiler–Leman procedure can distinguish between $\mathcal{A}(F)$ and $\mathcal{B}(F)$ within r refinement steps.

Proof sketch. Let us start by briefly recalling known characterizations in terms of Ehrenfeucht-Fraïssé games of L^k [Bar77, Imm82] and C^k [CFI92, Hel96]. In both cases the game is played by two players, called Spoiler and Duplicator, on the two structures $\mathcal{A}$ and $\mathcal{B}$. Positions in the games are partial mappings $p = \{(u_1, v_1), \dots, (u_i, v_i)\}$ from $V(\mathcal{A})$ to $V(\mathcal{B})$ of size at most k. The games start from the empty position and proceed in rounds. At the beginning of each round in both games, Spoiler chooses $p' \subseteq p$ with $|p'| < k$.

- In the L^k-game, Spoiler then selects either some $u \in V(\mathcal{A})$ or some $v \in V(\mathcal{B})$ and Duplicator responds by choosing an element $v \in V(\mathcal{B})$ or $u \in V(\mathcal{A})$ in the other structure.
- In the C^k-game, Duplicator first selects a global bijection $f : V(\mathcal{A}) \to V(\mathcal{B})$ and Spoiler chooses some pair $(u, v) \in f$. (If $|V(\mathcal{A})| \neq |V(\mathcal{B})|$, Spoiler wins the C^k-game immediately.)

The new position is $p' \cup \{(u, v)\}$. Spoiler wins the r-round L^k / C^k game if he has a strategy to reach within r rounds a position p that does not define an isomorphism on the induced substructures. Both games characterize equivalence in the corresponding logics: Spoiler wins the r-round L^k / C^k game if and only if there is a sentence $\varphi \in \mathsf{L}^k$ / C^k of quantifier depth r such that $\mathcal{A} \models \varphi$ and $\mathcal{B} \not\models \varphi$.

When these games are played on the two structures $\mathcal{A}(F)$ and $\mathcal{B}(F)$ obtained from an XOR formula F, it is not hard to verify that both games are equivalent to the k-pebble game on F. To see this, we identify Spoiler with Player 1, Duplicator with Player 2, and partial mappings $p = \{(x_i^{a_i}, x_i^{b_i}) \mid i \leq \ell\}$ with partial assignments $\alpha = \{x_i \mapsto a_i \oplus b_i \mid i \leq \ell\}$. Because of the X_i-relations, we can assume that partial assignments of any other form will not occur as they are losing positions for Duplicator.

If Spoiler asks for some x_i^0 or x_i^1 in the L^k-game, which corresponds to a choice by Player 1 of $x_i \in \mathit{Vars}(F)$, the only meaningful action for Duplicator is to choose either x_i^0 or x_i^1 in the other structure, corresponding to an assignment to x_i by Player 2. With any other choice Duplicator would lose immediately because of the unary relations X_i. Thus, there is a natural correspondence between strategies in the L^k-game and the k-pebble game.

The players in the k-pebble game can be assumed to have perfect knowledge of the strategy of the other player. This means that at any given position in the game, without loss of generality we can think of Player 1 as being given a complete truth value assignment to the remaining variables, out of which he can pick one variable assignment. By the correspondence in (2.2a)–(2.2b) we see that this can be translated to a bijection f chosen by Duplicator in the C^k-game (which has to preserve the X_i relations). Therefore, Spoiler picking some pair of the form (x_i^a, x_i^b) from f can be viewed as Player 1 asking about the assignment to x_i and getting a response from Player 2 in the game on F (again using the above-mentioned correspondence between partial mappings p and partial assignments α). Finally, we observe that by design a partial mapping that preserves the X_i-relations defines a local isomorphism if and only if the corresponding α does not falsify any XOR clause.

Formalizing the proof sketch above, it is not hard to show that statements (a)–(c) in the lemma are all equivalent. The equivalence between (c) and (d) was proven in [CFI92]. The lemma follows. □

3 Technical Lemmas and Proofs of Main Theorems

To prove our lower bounds of the quantifier depth of finite variable logics in Theorem 1.1 and the number of refinement steps of the Weisfeiler–Leman algorithm in Theorems 1.2 and 1.3, we utilize the characterization in Lemma 2.1 and show that there are n-variable XOR formulas on which Player 1 is able to win the k-pebble game but cannot do so in significantly less than $n^{k/\log k}$ rounds. The next lemma states this formally and also provides a trade-off as the number of pebbles increases.

Lemma 3.1 (Main technical lemma). *There is an absolute constant $K_0 \in \mathbb{N}^+$ such that for integers k_{lo}, k_{hi}, and n satisfying $K_0 \leq k_{\mathrm{lo}} \leq k_{\mathrm{hi}} \leq n^{1/6}/k_{\mathrm{lo}}$ there is an XOR formula F with n variables such that Player 1 wins the k_{lo}-pebble game on F, but does not win the k_{hi}-pebble game within $n^{k_{\mathrm{lo}}/(10\log k_{\mathrm{hi}})-1/5}$ rounds.*

Note that there is a limit to how far k_{lo} and k_{hi} can be from each other for the lemma to make sense—the statement becomes vacuous if $k_{\mathrm{lo}} \leq 2\log k_{\mathrm{hi}}$. Let us see how this lemma yields the theorems in Section 1.

Proof of Theorem 1.1. This theorem follows immediately from Lemmas 2.1 and 3.1, but let us write out the details for clarity. By setting $k_{\mathrm{lo}} = k_{\mathrm{hi}} = k$ in Lemma 3.1, we can find XOR formulas with n variables such that Player 1 wins the k-pebble game on F_n but needs more than $n^{\varepsilon k/\log k}$ rounds in order to do so (provided we choose $\varepsilon < 1/10$ and K_0 large enough). We can then plug these XOR formulas into Lemma 2.1 to obtain n-element structures $\mathcal{A}_n = \mathcal{A}(F_n)$ and $\mathcal{B}_n = \mathcal{B}(F_n)$ that can be distinguished in the k-variable fragments of first-order logic L^k and first-order counting logic C^k, but where this requires sentences of quantifier depth at least $n^{\varepsilon k/\log k}$. □

Proof of Theorem 1.2. If we let F_n be the XOR formula from Lemma 3.1 for $k_{\mathrm{lo}} = k_{\mathrm{hi}} = k+1$, then by Lemma 2.1 it holds that the structures $\mathcal{A}(F_n)$ and $\mathcal{B}(F_n)$ will be distinguished by the k-dimensional Weisfeiler–Leman algorithm, but only after $n^{\varepsilon(k+1)/\log(k+1)} \geq n^{\varepsilon k/\log k}$ refinement steps. Hence, computing the stable colouring of either of these structures requires at least $n^{\varepsilon k/\log k}$ refinement steps (since they would be distinguished earlier if at least one of the computations terminated earlier). □

Proof of Theorem 1.3. This is similar to the proof of Theorem 1.2, but setting $k_{\mathrm{lo}} = \lfloor \log n^c \rfloor + 1$ and $k_{\mathrm{hi}} = \lceil n^{1/7} \rceil + 1$ in Lemma 3.1. □

The proof of the trade-off between the number of pebbles versus number of rounds in Lemma 3.1 splits into two steps. We first establish a rather weak lower bound on the number of rounds in the pebble game played on suitably chosen m-variable XOR formulas for $m \gg n$. We then transform this into a much stronger lower bound for formulas over n variables using hardness condensation. To help the reader keep track of which results are proven in which setting, in what follows we will write ℓ_{lo} and ℓ_{hi} to denote parameters depending on m and k_{lo} and k_{hi} to denote parameters depending on n.

To implement the first step in our proof plan, we use tools developed by Immerman [Imm81] to establish a lower bound as stated in the next lemma.

Lemma 3.2. *For all $\ell_{\mathrm{hi}}, m \geq 3$ there is an m-variable 3-XOR formula $F_m^{\ell_{\mathrm{hi}}}$ on which Player 1*

(a) wins the 3-pebble game, but

(b) does not win the ℓ_{hi}-pebble game within $\max\bigl(3, \frac{1}{\lceil \log \ell_{\mathrm{hi}} \rceil} m^{1/(1+\lceil \log \ell_{\mathrm{hi}} \rceil)} - 2\bigr)$ rounds.

We defer the proof of Lemma 3.2 to Section 4, but at this point an expert reader might wonder why we would need to prove this lower bound at all, since a much stronger $\Omega(m)$ bound on the number of rounds in the pebble game on 4-XOR formulas was already obtained by Fürer [Fü01]. The reason is that in Fürer's construction Player 1 cannot win the game with few pebbles. However, it is crucial for the second step of our proof, where we boost the lower bound but also significantly increase the number of pebbles that are needed to win the game, that Player 1 is able to win the original game with very few pebbles.

The second step in the proof of our main technical lemma is carried out by using the techniques developed by Razborov [Raz16] and applying them to the XOR formulas in Lemma 3.2. Roughly speaking, if we set $k_{\mathrm{lo}} = k_{\mathrm{hi}} = k$ for simplicity, then the number of variables decreases from m to $n \approx m^{1/k}$, whereas the $m^{1/\log k}$ round lower bound for the k-pebble game stays essentially the same and hence becomes $n^{k/\log k}$ in terms of the new number of variables n. The properties of hardness condensation are summarized in the next lemma, which we prove in Section 5. To demonstrate the flexibility of this tool we state the lemma in its most general form—readers who want to see an example of how to apply it to the XOR formulas in Lemma 3.2 can mentally fix $p = 3$, $\ell_{\mathrm{lo}} = 3$, $\ell_{\mathrm{hi}} = k_{\mathrm{hi}}$, $r \approx m^{1/\log k_{\mathrm{hi}}}$, and $\Delta \approx k_{\mathrm{hi}}/3$ when reading the statement of the lemma below.

Lemma 3.3 (Hardness condensation lemma). *There exists an absolute constant $\Delta_0 \in \mathbb{N}^+$ such that the following holds. Let F be an m-variable p-XOR formula and suppose that we can choose parameters $\ell_{\mathrm{lo}} > 0$, $\ell_{\mathrm{hi}} \geq \Delta_0 \ell_{\mathrm{lo}}$ and r such that Player 1*

(a) has a winning strategy for the ℓ_{lo}-pebble game on F, but

(b) does not win the ℓ_{hi}-pebble game on F within r rounds.

Then for any Δ satisfying $\Delta_0 \leq \Delta \leq \ell_{\mathrm{hi}}/\ell_{\mathrm{lo}}$ and $(2\ell_{\mathrm{hi}}\Delta)^{2\Delta} \leq m$ there is an (Δp)-XOR formula H with $\lceil m^{3/\Delta} \rceil$ variables such that Player 1

(a') has a winning strategy for the $(\Delta \ell_{\mathrm{lo}})$-pebble game on H, but

(b') does not win the ℓ_{hi}-pebble game on H within $r/(2\ell_{\mathrm{hi}})$ rounds.

Taking Lemmas 3.2 and 3.3 on faith for now, we are ready to prove our main technical lemma yielding an $n^{\Omega(k/\log k)}$ lower bound on the number of rounds in the k-pebble game.

Proof of Lemma 3.1. Let Δ_0 be the constant in Lemma 3.3. We let

$$K_0 \geq 3\Delta_0 + 9 \tag{3.1}$$

be an absolute constant, the precise value of which will be determined by calculations later in the proof. We are given k_{hi}, k_{lo}, and n satisfying the conditions

$$K_0 \leq k_{\mathrm{lo}} \leq k_{\mathrm{hi}} \leq n^{1/6}/k_{\mathrm{lo}} \tag{3.2}$$

in Lemma 3.1. Let us set

$$\ell_{\mathrm{hi}} := k_{\mathrm{hi}} \tag{3.3a}$$

and

$$m := n^{\lfloor k_{\mathrm{lo}}/9 \rfloor} \tag{3.3b}$$

and apply Lemma 3.2 (which is in order since $\ell_{\mathrm{lo}} \geq 3$ and $m \geq 3$ by (3.1) and (3.2)). This yields an m-variable 3-XOR formula on which Player 1 wins the 3-pebble game but cannot win the ℓ_{hi}-pebble game within

$$r := \tfrac{1}{\lceil \log \ell_{\mathrm{hi}} \rceil} m^{1/(1+\lceil \log \ell_{\mathrm{hi}} \rceil)} - 2 \tag{3.3c}$$

rounds. As a side remark we note that this lower bound term might vanish if k_{lo} and k_{hi} were to far apart from each other ($k_{\mathrm{lo}} \leq 2 \log k_{\mathrm{hi}}$), but recall that in this case the statement of Lemma 3.1 becomes vacuous anyway. Now we can apply hardness condensation as in Lemma 3.3 to the formula provided by Lemma 3.2, where we fix parameters

$$p := 3 \ , \tag{3.4a}$$

$$\ell_{\mathrm{lo}} := 3 \ , \tag{3.4b}$$

and

$$\Delta := 3\lfloor k_{\text{lo}}/9 \rfloor \ . \tag{3.4c}$$

To verify that our choice of parameters is legal, note that in addition to $r \geq 1$ we also have $\ell_{\text{lo}} > 0$ and

$$\ell_{\text{hi}} = k_{\text{hi}} \geq K_0 > 3\Delta_0 = \Delta_0 \ell_{\text{lo}} \ . \tag{3.5}$$

Thus, the assumptions needed for (a) and (b) are satisfied by the XOR formula obtained from Lemma 3.2. To confirm that Δ chosen as in (3.4c) satisfies the conditions in Lemma 3.3, observe that

$$\Delta_0 \leq 3\lfloor K_0/9 \rfloor \leq 3\lfloor k_{\text{lo}}/9 \rfloor = \Delta \leq k_{\text{lo}}/3 \leq k_{\text{hi}}/3 = \ell_{\text{hi}}/\ell_{\text{lo}} \ . \tag{3.6}$$

Furthermore, since $\Delta \leq k_{\text{lo}}/3$ and $\ell_{\text{hi}} = k_{\text{hi}} \leq n^{1/6}/k_{\text{lo}}$ we get

$$(2\ell_{\text{hi}}\Delta)^{2\Delta} \leq \left(\frac{2}{3}n^{1/6}\right)^{2\Delta} \leq n^{\Delta/3} = m \ . \tag{3.7}$$

Note, finally, that $n = n^{\lfloor k_{\text{lo}}/9 \rfloor 3/\Delta} = m^{3/\Delta}$. Now Lemma 3.3 provides us with an n-variable k_{lo}-XOR formula on which according to (a') Player 1 has a winning strategy for the (3Δ)-pebble game and hence also for the game with $k_{\text{lo}} \geq 9\lfloor k_{\text{lo}}/9 \rfloor = 3\Delta$ pebbles. Moreover, by (b') it holds that Player 1 needs more than $r/(2k_{\text{hi}})$ rounds to win the k_{hi}-pebble game. To complete the proof, we observe that if we choose K_0 large enough, then for $n > k_{\text{hi}} \geq k_{\text{lo}} \geq K_0$ it holds that

$$\begin{aligned}
\frac{r}{2k_{\text{hi}}} &= \frac{1}{2k_{\text{hi}}\lceil \log k_{\text{hi}} \rceil} n^{\lfloor k_{\text{lo}}/9 \rfloor/(1+\lceil \log k_{\text{hi}} \rceil)} - \frac{1}{k_{\text{hi}}} && \big[\text{by (3.3b) and (3.3c)}\big] \\
&\geq \frac{6n^{1/5}}{n^{1/6}\log n} n^{\lfloor k_{\text{lo}}/9 \rfloor/(1+\lceil \log k_{\text{hi}} \rceil)-1/5} - \frac{1}{k_{\text{hi}}} && \big[\text{since } k_{\text{hi}} \leq n^{1/6}\big] \\
&\geq n^{k_{\text{lo}}/(10 \log k_{\text{hi}})-1/5} && \big[\text{for large enough } n, k_{\text{hi}}, \text{ and } k_{\text{lo}}.\big]
\end{aligned} \tag{3.8}$$

We now choose the constant K_0 large so that all conditions encountered in the calculations above are valid. This establishes the lemma. □

4 XOR Formulas over High-Dimensional Pyramids

We now proceed to establish the k-pebble game lower bound stated in Lemma 3.2. Our XOR formulas will be constructed over directed acyclic graphs (DAGs) as described in the following definition.

Definition 4.1. Let $\mathcal{G}$ be a DAG with sources S and a unique sink z. The XOR formula $xor(\mathcal{G})$ contains one variable v for every vertex $v \in V(\mathcal{G})$ and consists of the following clauses:

(a) $(s, 0)$ for every source $s \in S$,

(b) $(v, w_1, \ldots, w_\ell, 0)$ for all non-sources $v \in V(\mathcal{G}) \setminus S$ with in-neighbours $N^-(v) = \{w_1, \ldots, w_\ell\}$,

(c) $(z, 1)$ for the unique sink z.

Note that the formula $xor(\mathcal{G})$ is always unsatisfiable, since all source vertices are forced to 0 by (a), which forces all other vertices to 0 in topological order by (b), contradicting (c) for the sink. Incidentally, these formulas are somewhat similar to the *pebbling formulas* defined in [BW01], which have been very useful in proof complexity (see the survey [Nor13] for more details). The difference is that pebbling formulas state that a vertex v is true if and only if all of its in-neighbours are true, whereas $xor(\mathcal{G})$ states that v is true if and only if the parity of the number of true in-neighbours is odd.

It is clear that one winning strategy for Player 1 is to ask first about the sink z, for which Player 2 has to answer 1 (or lose immediately) and then about all the in-neighbours of the sink until the answer for

one vertex v is 1 (if there is no such vertex, Player 2 again loses immediately). At this point Player 1 can forget all other vertices and then ask about the in-neighbours of v until a 1-labelled vertex w is found, and then continue in this way to trace a path of 1-labelled vertices backwards through the DAG until some source s is reached, which contradicts the requirement that s should be labelled 0. Formalizing this as an induction proof on the depth of $\mathcal{G}$ shows that if the in-degree is bounded, then Player 1 can win the pebble game on $xor(\mathcal{G})$ with few pebbles as stated in the next lemma.

Lemma 4.2. *Let $\mathcal{G}$ be a DAG with a unique sink and maximal in-degree d. Then Player 1 wins the $(d+1)$-pebble game on $xor(\mathcal{G})$.*

As a warm-up for the proof of Lemma 3.2, let us describe a very weak lower bound from [Imm81] for the complete binary tree of height h (with edges directed from the leaves to the root), which we will denote $\mathcal{T}_h$. By the lemma above, Player 1 wins the 3-pebble game on $xor(\mathcal{T}_h)$ in $O(h)$ steps by propagating 1 from the root down to some leaf. On the other hand, Player 2 has the freedom to decide on which path she answers 1. Hence, she can safely respond 0 for a vertex v as long as there is some leaf with a pebble-free path leading to the lowest pebble labelled 1 without passing v. In particular, if Player 2 is asked about vertices at least ℓ layers below the lowest pebbled vertex for which the answer 1 was given, then she can answer 0 for $2^\ell - 1$ queries. It follows that the height h provides a lower bound on the number of rounds Player 1 needs to win the game, even if he has an infinite amount of pebbles. We remark that this proof in terms of pebble-free paths is somewhat reminiscent of an argument by Cook [Coo74] for the so-called black pebble game corresponding to the pebbling formulas in [BW01] briefly discussed above.

The downside of this lower bound is that the height is only logarithmic in the number of vertices and thus too weak for us as we are shooting for a lower bound of the order of $n^{1/\log k}$. To get a better bound for the black pebble game Cook instead considered so-called pyramid graphs as in Figure 2(a). These will not be sufficient to obtain strong enough lower bounds for our pebble game, however.[2] Instead, following Immerman we consider a kind of high-dimensional generalization of these graphs, for which the lower bound on the number of rounds in the k-pebble game is still linear in the height h while the number of vertices is roughly $h^{\log k}$.

Definition 4.3 ([Imm81]). For $d \geq 1$ we define the $(d+1)$-*dimensional pyramid of height* h, denoted by $\mathcal{P}_h^d$, to be the following layered DAG. We let L, $0 \leq L \leq h$ be the *layer number* and set $q_d(L) := \lfloor L/d \rfloor$ and $r_d(L) := L \pmod d$. Hence, for any L we have $L = q_d(L) \cdot d + r_d(L)$. For integers $x_i \geq 0$ the vertex set is

$$V\big(\mathcal{P}_h^d\big) = \big\{(x_0, \ldots, x_{d-1}, L) \;\big|\; L \leq h;\; x_i \leq q_d(L)+1 \text{ if } i < r_d(L);\; x_i \leq q_d(L) \text{ if } i \geq r_d(L)\big\} \;, \tag{4.1a}$$

where we say that L is the *layer* of the vertex $(x_0, \ldots, x_{d-1}, L)$. The edge set $E\big(\mathcal{P}_h^d\big)$ consists of the pair of edges

$$\begin{gathered}\big((x_0, \ldots, x_{r_d(L)}, \ldots, x_{d-1}, L+1), (x_0, \ldots, x_{r_d(L)}, \ldots, x_{d-1}, L)\big) \;, \\ \big((x_0, \ldots, x_{r_d(L)}+1, \ldots, x_{d-1}, L+1), (x_0, \ldots, x_{r_d(L)}, \ldots, x_{d-1}, L)\big)\end{gathered} \tag{4.1b}$$

for all vertices $(x_0, \ldots, x_{d-1}, L) \in V(\mathcal{P}_h^d)$ and layers $L < h$, so that every vertex in layer L has exactly two in-neighbours from layer $L+1$.

It might be easier to parse Definition 4.3 by noting that the (kd)th layer of $\mathcal{P}_h^d$ is a d-dimensional cube of side length k. Intuitively, we then want to have incoming edges to each vertex u at the (kd)th

[2] For readers knowledgeable in pebbling, we comment that the problem is that the open-path argument in [Coo74] does not work in a DAG-like setting for the XOR pebble game. To see this, consider a pyramid with a vertex row u, v, w and a second row p, q, r, s immediately below such that the edges are $(p,u), (q,u), (q,v), (r,v), (r,w), (s,w)$. Then if the values of u, w on the upper row and p, s on the lower row are known, there is still an open path via (q,v) or (r,v), which is enough for the black pebbling lower bound for pyramids in [Coo74]. But in the XOR pebble game this means that r and q are already fixed because of the XOR constraints, and so there is no "open path" with unconstrained vertices.

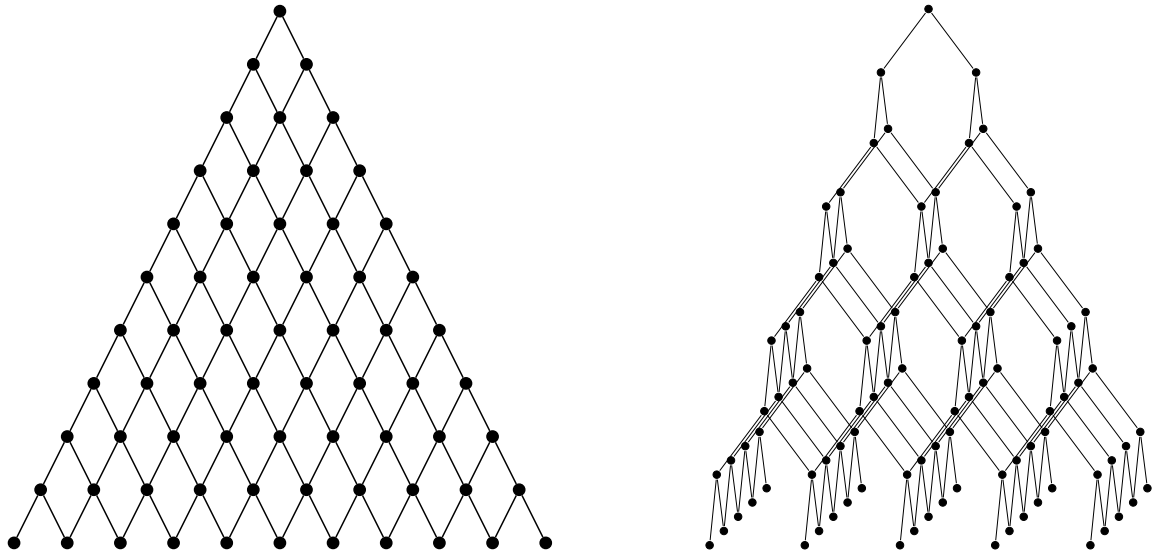

(a) 2D pyramid (b) 3D pyramid

Figure 2: Examples of high-dimensional pyramids.

layer from all vertices v in the d-dimensional cube of side length $k+1$ such that all coordinates of v are at distance 0 or $+1$ from the coordinates of u. This would give a fan-in larger than 2, however, and to avoid this we expand in one dimension at a time to obtain a sequence of multidimensional cuboids where in each consecutive cuboid the side length increases by one in one dimension, until d layers later we have a complete cube with side length $k+1$. We refer the reader to Figure 2(a) for an illustration of a 2-dimensional pyramid generated by stacking 1-dimensional cubes on top of one another and to Figure 2(b) for a 3-dimensional pyramid generated from 2-dimensional cuboids (where all the edges in the figures are assumed to be directed upwards). The vertex $(0,\ldots,0)$ at the top of the pyramid is the unique sink and all vertices at the bottom layer h are sources. Observe that it follows from the definition that $\left|V\left(\mathcal{P}_h^d\right)\right| \leq (h+1)^{d+1}$.

As high-dimensional pyramids have in-degree 2, Lemma 4.2 implies that Player 1 wins the 3-pebble game on $\mathcal{P}_h^d$. Recall that, as discussed in the proof sketch of the lemma, Player 1 starts his winning strategy in the 3-pebble game by pebbling the sink of the pyramid and its two in-neighbours. One of them has to be labelled 1. Then he picks up the two other pebbles and pebbles the two in-neighbours of the vertex marked with 1 and so on. Continuing this strategy, he is able to "move" the 1 all the way to the bottom, reaching a contradiction, in a number of rounds that is linear in the height of the pyramid. This strategy turns out to be nearly optimal in the sense that in order to move a 1 from the top to the bottom in $\mathcal{P}_h^d$, as long as the total number of available pebbles is at most 2^d it makes no sense for Player 1 at any point in the game to pebble a vertex that is d or more levels away from the lowest level containing a pebble.

The next lemma states a key property of pyramids in this regard. In order to state it, we need to make a definition.

Definition 4.4. We refer to a partial assignment $\mathcal{M}$ of Boolean values to the vertices of a DAG $\mathcal{G}$ as a *labelling* or *marking* of $\mathcal{G}$. We say that $\mathcal{M}$ is *consistent* if no clause of type (b) or (c) in the XOR formula $xor(\mathcal{G})$ in Definition 4.1 is falsified by $\mathcal{M}$. We also say that $\mathcal{M}'$ is *consistent with* $\mathcal{M}$ if $\mathcal{M} \cup \mathcal{M}'$ is a consistent labelling of $\mathcal{G}$.

That is, a consistent labelling does not violate any constraint on any non-source vertex, but source vertex constraints (a) may be falsified. Such labellings are easy to find for high-dimensional pyramids.

Lemma 4.5 ([Imm81]). *Let $\mathcal{M}$ be any consistent labelling of all vertices in a pyramid $\mathcal{P}_h^d$ from layer 0 to layer L. Then for every set S of 2^d-1 vertices on or below layer $L+d$ there is a consistent labelling*

of the entire pyramid that extends $\mathcal{M}$ and labels all vertices in S with 0.

To get some intuition why Lemma 4.5 holds, note that the d-dimensional pyramids are constructed in such a way that they locally look like binary trees. In particular, every vertex $v \in V(\mathcal{P}_h^d)$ together with all its predecessors at distance at most d form a complete binary tree. By the same argument as for the binary trees above, it follows that if v is labelled with 1, Player 2 can safely answer 0 up to $2^d - 1$ times when asked about vertices d layers below v. However, the full proof of Lemma 4.5 is more challenging and requires some quite subtle reasoning. For the convenience of the reader we now present a slightly modified version of the proof in [Imm81] with notation and terminology adapted to this paper.

To formalize the intuitive argument above, we need some additional notation and technical definitions. Let us use the shorthand $\vec{x} = (x_0, \ldots, x_{d-1})$. For a pyramid $\mathcal{P}_h^d$, a coordinate $j \in \{0, \ldots, d-1\}$, and a layer L, we let

$$\mathit{slen}(j, L) := \max\{x_j \mid (\vec{x}, L) \in V(\mathcal{P}_h^d)\} \tag{4.2}$$

be the side length in the jth dimension of the cuboid in layer L, i.e., the maximal value that can be achieved in the jth coordinate in layer L, and for $L' \geq L$ we write

$$\Delta_{\mathit{slen}}(j, L, L') := \mathit{slen}(j, L') - \mathit{slen}(j, L) \tag{4.3}$$

to denote how much the cuboids in $\mathcal{P}_h^d$ grow in the jth dimension in betwen layers L and L'.

We define the *frustum* $\mathcal{P}_{L,h}^d$ to be the subgraph of $\mathcal{P}_h^d$ induced on the set $\{(\vec{x}, L') \mid L' \geq L\}$ of all vertices on layer L and below. We say that the *wedge* $\mathcal{W}(j, a, L)$ is the subgraph of $\mathcal{P}_h^d$ induced on the vertices $(x_0, \ldots, x_{j-1}, a, x_{j+1}, \ldots, x_{d-1}, L)$ with fixed jth coordinate $x_j = a$ together with all predecessors of these vertices. That is, the vertex set of $\mathcal{W}(j, a, L)$ is

$$V(\mathcal{W}(j, a, L)) = \{(\vec{x}, L') \in V(\mathcal{P}_h^d) \mid L' \geq L, a \leq x_j \leq a + \Delta_{\mathit{slen}}(j, L, L')\} \ . \tag{4.4}$$

An important part in our proof will be played by subgraphs obtained by deleting wedges from frustums. We define these subgraphs next.

Fix a frustum $\mathcal{P}_{L,h}^d$, two disjoint subsets of coordinates $I_{\text{lo}}, I_{\text{hi}} \subseteq \{0, \ldots, d-1\}$, and a mapping $\alpha \colon I_{\text{lo}} \cup I_{\text{hi}} \to \mathbb{N}_0$ such that $\alpha(j) \leq \mathit{slen}(j, L)$. We let the *restricted frustum* $\mathcal{P}_{L,h}^d[I_{\text{lo}}, I_{\text{hi}}, \alpha]$ be the subgraph of the frustum $\mathcal{P}_{L,h}^d$ induced on the vertex set

$$\{(\vec{x}, L') \in V(\mathcal{P}_{L,h}^d) \mid \forall j \in I_{\text{lo}} : x_j > \alpha(j) + \Delta_{\mathit{slen}}(j, L, L'); \forall j \in I_{\text{hi}} : x_j < \alpha(j)\} \tag{4.5}$$

where no coordinates in $I_{\text{lo}} \cup I_{\text{hi}}$ are expanded, i.e., the cuboids will not grow in size in dimensions $I_{\text{lo}} \cup I_{\text{hi}}$ as we move down the layers. For dimensions in I_{hi} the coordinate set stays the same, and for dimensions in I_{lo} the coordinate set shifts by an additive $+1$ every time the pyramid graph grows in this direction. We say that a layered directed graph is a (d, q)-*frustum* if it is a restricted frustum $\mathcal{P}_{L,h}^d[I_{\text{lo}}, I_{\text{hi}}, \alpha]$ where q coordinates are restricted, i.e., $|I_{\text{lo}} \dot\cup I_{\text{hi}}| = q$.

To see how restricted frustums are obtained by deleting wedges from frustums, note that after removing the wedge $\mathcal{W}(j, a, L)$ from the frustum $\mathcal{P}_{L,h}^d$, the remaining graph is the disjoint union of the restricted frustums $\mathcal{P}_{L,h}^d[\{j\}, \emptyset, \{j \mapsto a\}]$ and $\mathcal{P}_{L,h}^d[\emptyset, \{j\}, \{j \mapsto a\}]$. Figure 3 shows a 3D pyramid with a wedge ○ and a restricted frustum ✕.

We prove Lemma 4.5 by inductively cutting the pyramid into a wedge and restricted frustums to the left and right of this wedge. It will be convenient to focus on (d, q)-frustums which grow in the dimensions corresponding to the topmost $d - q$ layers (as the one in Figure 3). More formally, we say that an (d, q)-frustum $\mathcal{P}_{L,h}^d[I_{\text{lo}}, I_{\text{hi}}, \alpha]$ is *top-expanding* if

$$\{(L + j) \bmod d \mid 0 \leq j \leq d - 1 - q\} \cap (I_{\text{lo}} \cup I_{\text{hi}}) = \emptyset \ . \tag{4.6}$$

As we have done for digraphs with unique sinks in Definition 4.1, we identify with each (restricted) frustum $\mathcal{P}$ the corresponding XOR formula $\mathit{xor}(\mathcal{P})$ containing all clauses given in Definition 4.1(a)

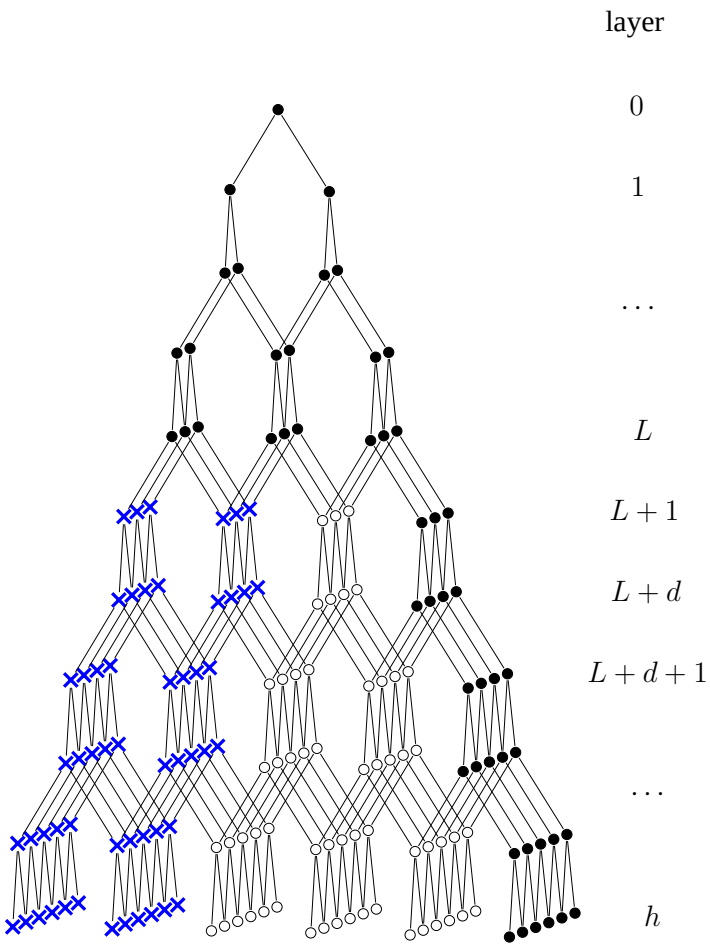

Figure 3: Pyramid with wedge $\mathcal{W}(0, 2, L+1)$ ○ and restricted frustum $\mathcal{P}^2_{L+1,10}[\emptyset, \{0\}, \{0 \mapsto 2\}]$ ✕.

and (b). We do not include hard-coded labels on the sources at the top layer as in (c) but instead will always provide a labelling of that layer.[3]

We now state our inductive claim. Lemma 4.5 follows immediately once this claim has been established, as the subgraph of a pyramid $\mathcal{P}^d_h$ on or below layer L is an (unrestricted) top-expanding $(d, 0)$-frustum $\mathcal{P}^d_{L,h}$ and, in particular, $xor(\mathcal{P}^d_h)$ with all vertices on or above layer L consistently labelled is equivalent to $xor(\mathcal{P}^d_{L,h})$ with the same labelling of layer L.

Claim 4.6. Let $\mathcal{P}$ be a top-expanding (d, q)-frustum and $\mathcal{M}_L$ be a labelling of its top layer L. Then for every set S of $2^{d-q} - 1$ vertices on or below layer $L + d - q$ there is a consistent labelling of all vertices in $\mathcal{P}$ that extends $\mathcal{M}_L$ and labels every vertex in S with 0.

The following proposition summarizes the core properties of frustums that we will use when establishing Claim 4.6.

Proposition 4.7. *Let $\mathcal{P} = \mathcal{P}^d_{L,h}[I_{\text{lo}}, I_{\text{hi}}, \alpha]$ be a restricted frustum with a labelling $\mathcal{M}_L$ of all vertices in the top layer L. Then the following holds:*

(a) There is a labelling $\mathcal{M}_{L+1}$ of all vertices in layer $L+1$ of $\mathcal{P}$ that is consistent with $\mathcal{M}_L$.

(b) Let $\mathcal{M}_{L+1}$ be any labelling of layer $L+1$ in $\mathcal{P}$ that is consistent with $\mathcal{M}_L$ and suppose that $\mathcal{P}$ expands coordinate j from layer L to layer $L+1$, i.e., $j = r_d(L)$ and $j \notin I_{\text{lo}} \cup I_{\text{hi}}$. Then for any

[3] For readers more familiar with proof complexity language, our subgraphs correspond to formulas obtained by applying restrictions to $xor(\mathcal{P}^d_h)$.

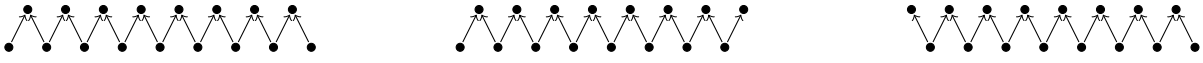

(a) $j \notin I_{\text{lo}} \cup I_{\text{hi}}$ (b) $j \in I_{\text{hi}}$ (c) $j \in I_{\text{lo}}$

Figure 4: Shapes of connected component between layers L and $L+1$ expanding in dimension j.

vertex $(\vec{y}, L+1)$ *in* $\mathcal{P}$ *it holds that the labelling* $\mathcal{M}^{\vec{y}}_{L+1}$ *defined by*

$$\mathcal{M}^{\vec{y}}_{L+1}(\vec{x}, L+1) := \begin{cases} 1 - \mathcal{M}_{L+1}(\vec{x}, L+1) & \text{if } x_i = y_i \text{ for all } i \neq j, \\ \mathcal{M}_{L+1}(\vec{x}, L+1) & \text{otherwise} \end{cases}$$

is also consistent with $\mathcal{M}_L$.

Proof. We first note that the set of XOR constraints between two layers L and $L+1$ can be partitioned into several connected components. Each of the components forms a "one-dimensional line" in the direction of the expanding coordinate $j = r_d(L)$. More formally, two vertices from layers L and $L+1$ are on the same line if and only if they agree on the coordinates x_i for all $i \neq j$. These lines form the connected components of the graph induces on layers L and $L+1$. All such lines between layers L and $L+1$ isomorphic and their shape depends on whether $j \in I_{\text{lo}}$, $j \in I_{\text{hi}}$, or $j \notin I_{\text{lo}} \cup I_{\text{hi}}$. See Figure 4 for an illustration. We remark that in Figures 4(b) and 4(c) the vertex with in-degree 1 and its predecessor form a binary XOR clause in $xor(\mathcal{P})$.

If the layer is not expanding (as depicted in Figures 4(b) and 4(c)) and the upper layer L is entirely labelled, then it is not hard to see that there is a unique consistent labelling of the lower-level vertices of each line (determined by propagating values from right to left in Figure 4(b) and from left to right in Figure 4(c)). As all lines are disjoint this gives a unique labelling of the entire layer $L+1$ that is consistent with the labelling of layer L. If the layer expands (i.e., , if $j \notin I_{\text{lo}} \cup I_{\text{hi}}$ as illustrated in Figure 4(a)), then we have more freedom. Indeed, if we label either the rightmost or the leftmost vertex at the bottom layer with 0, then we have the same situation as in Figures 4(b) and 4(c), respectively. This concludes the proof of item (a) in the proposition.

For item (b), first observe that the condition $j \notin I_{\text{lo}} \cup I_{\text{hi}}$ means that we are in the case depicted in Figure 4(a). This means that if we have a consistent labelling of the upper and lower part of a line, then flipping all values at the lower level yields another consistent labelling. This is so since every XOR clause contains exactly two vertices from the lower part. Hence, flipping both of these vertices does not change the parity of the variables in the XOR clause but leaves the clause satisfied. As all lines between layers L and $L+1$ are disconnected from each other, flipping all values in one line gives another consistent labelling for the whole layer $L+1$, which is precisely what is claimed in item (b). The proposition follows. □

Proof of Claim 4.6. The proof is by induction over decreasing values of q, the base case being $q = d$. As $|S| = 2^{d-q} - 1 = 0$ if $q = d$, in this case we only have to ensure that there is a consistent labelling of the entire frustum that is consistent with the labelling of the top layer. This follows from inductively applying Proposition 4.7(a) layer by layer.

For the inductive step, assume that the claim holds for all top-expanding $(d, q+1)$-frustums. We want prove it for a top-expanding (d, q)-frustum $\mathcal{P} = \mathcal{P}^d_{L,h}[I_{\text{lo}}, I_{\text{hi}}, \alpha]$. Let $j = r_d(L) = L \bmod d$ be the dimension that expands from layer L to $L+1$. As $\mathcal{P}$ is top-expanding and $q < d$ we have $j \notin I_{\text{lo}} \cup I_{\text{hi}}$. For some well-chosen $a \in [0, slen(j, L+1)]$ to be specified shortly, we partition $\mathcal{P}$ on and below layer $L+1$ into the wedge

$$\mathcal{W}(j, a, L+1) \tag{4.7a}$$

and two disjoint $(d, q+1)$-frustums: the "right" frustum

$$\mathcal{P}^d_{L+1,h}[I_{\text{lo}} \cup \{j\}, I_{\text{hi}}, \alpha \cup \{j \mapsto a\}] \tag{4.7b}$$

and the "left" frustum (depicted by ✕ in Figure 3)

$$\mathcal{P}^d_{L+1,h}[I_{\text{lo}}, I_{\text{hi}} \cup \{j\}, \alpha \cup \{j \mapsto a\}] \ . \tag{4.7c}$$

We choose the position a of the wedge so that both $(d, q+1)$-frustums in (4.7b) and (4.7c) contain at most $(|S|-1)/2 \leq 2^{d-(q+1)} - 1$ vertices from S (which implies that the wedge (4.7a) contains at least one vertex from S). To be more specific, we choose the largest $a \geq 0$ such that the left frustum (4.7c) contains at most $(|S|-1)/2$ vertices $S_a \subseteq S$. Such an a exists as $S_0 = \emptyset$.

If a reached the maximum $slen(j, L+1)$, then empty right frustum (4.7b) clearly contains no vertices from S. Otherwise let S_{a+1} be the set of vertices from S left of the wedge at position $a+1$. By the choice of a we have $|S_{a+1}| > (|S|-1)/2$. Furthermore, because all vertices in S are below layer $L+1$ it follows that $S_{a+1} \setminus S_a$ is contained in the wedge (4.7a) at position a. Hence the right frustum (4.7b) contains at most $|S \setminus S_{a+1}| \leq (|S|-1)/2$ vertices.

Now we proceed as follows. First we use Proposition 4.7(a) to obtain any consistent labelling of all vertices in layer $L+1$. Consider the set of all vertices $v = (\vec{x}, L+1)$ in layer $L+1$ with $x_j = a$, i.e., the topmost vertices in the wedge (4.7a). Note that this set of vertices form a hyperplane through, and perpendicular to, the disconnected parallel lines discussed in Proposition 4.7. We go over these vertices v one by one and flip every 1-labelled v to 0. As j is the expanding coordinate from layer L to $L+1$, after every such flip we can apply Proposition 4.7(b) to relabel the rest of the line through v as needed. In this way, all vertices in the top layer of the wedge (4.7a), get labelled by 0, and we label all other vertices in the wedge with 0 also. It follows from repeated application of Proposition 4.7(b) that the end result is a labelling of layer $L+1$ that is consistent with $\mathcal{M}_L$. Note that this labels every vertex from S within the wedge with 0 and moreover, layer $L+1$ contains no vertices from S outside of the wedge. This is because if some vertex from S is on layer $L+1$, then we have by the assumption in Claim 4.6 that $|S| = 1$ and hence this one labelled vertex is guaranteed to be in the wedge by the choice of a. In this way we obtain a labelling for the top layer of both $(d, q+1)$-frustums, and we then apply induction to consistently label all vertices in both frustums in such a way that every vertex in S is set to 0. Now we argue that the disjoint union of the all-zero labelling of the wedge and the consistent labellings of both frustums is a consistent labelling of $\mathcal{P}$. Clearly, every (non-source) clause in $xor(\mathcal{P})$ that is entirely contained in the wedge is satisfied by the all-zero labelling of the wedge. In the same way, every clause that is entirely contained in one of the two frustums is satisfied by their consistent labellings. It remains to consider clauses that contain variables from the wedge as well as from one of the frustums. By construction, those clauses have the form $(v, w_1, w_2, 0)$, for a vertex v with in-neighbours w_1, w_2, where the vertex v and one of its neighbours (say w_1) is within the frustum and the other neighbour is inside the wedge. As the edge (w_1, v) inside the frustum forms the binary clause $(v, w_1, 0)$ it follows that the consistent labelling of the frustum guarantees that the parity of v and w_1 is even. Because w_2 is labelled 0 in the wedge, the merged labelling satisfies $(v, w_1, w_2, 0)$. The claim follows. □

As noted above, our proof of Claim 4.6 also establishes Lemma 4.5.

In [Imm81] $\log n$-dimensional pyramids (where n is the number of vertices) are used to prove a $\Omega\big(2^{\sqrt{\log n}}\big)$ lower bound on the quantifier depth of full first-order counting logic. The next lemma shows that if we instead choose the dimension to be logarithmic in the number of variables (i.e., pebbles) in the game, we get an improved quantifier depth lower bound for the k-variable fragment.

Lemma 4.8. *For every $d \geq 2$ and height h, Player 1 does not win the 2^d-pebble game on $xor\big(\mathcal{P}^d_h\big)$ within $h/(d-1) - 1$ rounds.*

Proof. We show that Player 2 has a counter-strategy to answer consistently for at least $\lfloor h/(d-1) \rfloor - 1$ rounds and therefore Player 1 needs at least $\lfloor h/(d-1) \rfloor > h/(d-1) - 1$ rounds to win. Starting at the top layer $L_1 = 0$, she maintains the invariant that at the start of round r she has a consistent labelling of all vertices from layer 0 to layer L_r with the property that there is no pebble on layers L_r+1 to L_r+d-1.

Whenever Player 1 places a pebble on or above layer L_r, Player 2 responds according to the consistent labelling and whenever Player 1 puts a pebble on or below layer $L_r + d$, she answers 0, and in

both cases sets $L_{r+1} = L_r$. Note that as long as Player 1 places pebbles in this way, the game can go on forever. Since there are never more than $2^d - 1$ pebbles left on vertices on or below layer $L_r + d$ (when Player 1 runs out of pebbles the next move must be a removal), the conditions needed for Lemma 4.5 to apply are never violated.

Thus, the interesting case is when Player 1 places a pebble between layer $L_r + 1$ and $L_r + d - 1$. Then Player 2 uses Lemma 4.5 to extend her labelling to the first layer $L_{r+1} > L_r$ such that there is no pebble on layers $L_{r+1} + 1$ to $L_{r+1} + d - 1$, after which she answers the query according to the new labelling. It is worth noting that when Player 2 skips downward from layer L_r to layer L_{r+1} she might jump over a lot of layers in one go, but if so there is at least one pebble for every $(d-1)$th layer forcing such a big jump. We see that following this strategy Player 2 survives for at least $\lfloor h/(d-1) \rfloor - 1$ rounds, and this establishes the lemma. □

Putting the pieces together, we can now present the lower bound for the k-pebble game in Lemma 3.2.

Proof of Lemma 3.2. Recall that we want to prove that for all $\ell_{\text{hi}} \geq 3$ and $m \geq 3$ there is an m-variable 3-XOR formula F on which Player 1 wins the 3-pebble game but cannot win the ℓ_{hi}-pebble game within $\frac{1}{\lceil \log \ell_{\text{hi}} \rceil} m^{1/(1+\lceil \log \ell_{\text{hi}} \rceil)} - 2$ rounds. If $m < (5\lceil \log \ell_{\text{hi}} \rceil)^{(\lceil \log \ell_{\text{hi}} \rceil + 1)}$, then the round lower bound is trivial and we let F be, for instance, the 3-variable formula $xor\big(\mathcal{P}_1^1\big)$ plus $m - 3$ auxiliary variables on which Player 1 needs 3 rounds to win. Otherwise, we choose the formula to be $F = xor\big(\mathcal{P}_h^d\big)$ for parameters $d = \lceil \log \ell_{\text{hi}} \rceil$ and $h = \big\lfloor m^{1/(d+1)} \big\rfloor - 1$. Note that $\mathcal{P}_h^d$ contains less than $(h+1)^{d+1} \leq m$ vertices and we can add dummy variables to reach exactly m. Since the graph $\mathcal{P}_h^d$ has in-degree 2, Lemma 4.2 says that Player 1 wins the 3-pebble game as claimed in Lemma 3.2(a). The lower bound for the ℓ_{hi}-pebble game in Lemma 3.2(b) follows from Lemma 4.8 and the oberservation that because $h \geq 5d - 1$ we have $h/(d-1) \geq (h+1)/d$ and hence

$$h/(d-1) - 1 \leq (h+1)/d - 1 = \frac{1}{\lceil \log \ell_{\text{hi}} \rceil} \left\lfloor m^{1/(1+\lceil \log \ell_{\text{hi}} \rceil)} \right\rfloor - 1 \tag{4.8}$$

$$\leq \frac{1}{\lceil \log \ell_{\text{hi}} \rceil} m^{1/(1+\lceil \log \ell_{\text{hi}} \rceil)} - 2 \ . \tag{4.9}$$

The lemma follows. □

5 Hardness Condensation

In this section we establish Lemma 3.3, which shows how to convert an XOR formula into a harder formula over fewer variables. As discussed in the introduction, this part of our construction relies heavily on Razborov's recent paper [Raz16]. We follow his line of reasoning closely below, but translate it from proof complexity to a pebble game argument for bounded variable logics.

A key technical concept in the proof is graph expansion. Let us define the particular type of expander graphs we need and then discuss some crucial properties of these graphs. We use standard notation, letting $\mathcal{G} = (U \dot\cup V, E)$ denote a bipartite graph with left vertex set U and right vertex set V. We let $N^{\mathcal{G}}\big(U'\big) = \big\{v \,\big|\, \{u, v\} \in E(\mathcal{G}), u \in U'\big\}$ denote the set of neighbour vertices on the right of a left vertex subset $U' \subseteq U$ (and vice versa for right vertex subsets).

Definition 5.1 (Boundary expander). A bipartite graph $\mathcal{G} = (U \dot\cup V, E)$ is an $m \times n$ (s, c)-*boundary expander graph* if $|U| = m$, $|V| = n$, and for every set $U' \subseteq U$, $|U'| \leq s$, it holds that $\big|\partial^{\mathcal{G}}(U')\big| \geq c|U'|$, where the *boundary* $\partial^{\mathcal{G}}(U')$ is the set of all $v \in N^{\mathcal{G}}(U')$ having a unique neighbour in U', meaning that $\big|N^{\mathcal{G}}(v) \cap U'\big| = 1$. An (s, Δ, c)-boundary expander is an (s, c)-boundary expander where additionally $\big|N^{\mathcal{G}}(u)\big| \leq \Delta$ for all $u \in U$, i.e., the graph has left degree bounded by Δ.

In what follows, we will omit $\mathcal{G}$ from the notation when the graph is clear from context.

In any (s, c)-boundary expander with expansion $c > 0$ it holds that any left vertex subset $U' \subseteq U$ of size $|U'| \leq s$ has a partial matching into V where in addition the vertices in U' can be ordered in such a way that every vertex $u_i \in U'$ is matched to a vertex outside of the neighbourhood of the preceding vertices $u_1, \ldots, u_{i-1}$. The proof of this fact is sometimes referred to as a *peeling argument*.

Lemma 5.2 (Peeling lemma). *Let $\mathcal{G} = (U \dot\cup V, E)$ be an (s, c)-boundary expander with $s \geq 1$ and $c > 0$. Then for every set $U' \subseteq U$, $|U'| = t \leq s$ there is an ordering $u_1, \ldots, u_t$ of its vertices and a sequence of vertices $v_1, \ldots, v_t \in V$ such that $v_i \in N(u_i) \setminus N(\{u_1, \ldots, u_{i-1}\})$.*

Proof. The proof is by induction on t. The base case $t = 1$ is immediate since $s \geq 1$ and $c > 0$ implies that no left vertex can be isolated. For the inductive step, suppose the lemma holds for $t - 1$. To construct the sequence $v_1, \ldots, v_t$ we first fix v_t to be any vertex in $\partial(U')$, which has to exist since $\big|\partial(U')\big| \geq c|U'| > 0$. The fact that v_t is in the boundary of U' means that there is a unique $u_t \in U'$ such that $\big|N(v_t) \cap U'\big| = \{u_t\}$. Thus, for this pair (u_t, v_t) it holds that $v_t \in N(u_t) \setminus N\big(U' \setminus \{u_t\}\big)$. By the induction hypothesis we can now find sequences $u_1, \ldots, u_{t-1}$ and $v_1, \ldots, v_{t-1}$ for $U' \setminus \{u_t\}$ such that $v_i \in N(u_i) \setminus N(\{u_1, \ldots, u_{i-1}\})$, to which we can append u_t and v_t at the end. The lemma follows. □

For a right vertex subset $V' \subseteq V$ in $\mathcal{G} = (U \dot\cup V, E)$ we define the *kernel* $\mathrm{Ker}\big(V'\big) \subseteq U$ to be the set of all left vertices whose entire neighbourhood is contained in V', i.e.,

$$\mathrm{Ker}\big(V'\big) = \big\{u \in U \,\big|\, N(u) \subseteq V'\big\} \ . \tag{5.1}$$

We let $\mathcal{G} \setminus V'$ denote the subgraph of $\mathcal{G}$ induced on $\big(U \setminus \mathrm{Ker}(V')\big) \dot\cup \big(V \setminus V'\big)$. In other words, we obtain $\mathcal{G} \setminus V'$ from $\mathcal{G}$ by first deleting V' and afterwards all isolated vertices from U (assuming that there were no isolated left vertices before, which is true if $\mathcal{G}$ is expanding).

The next lemma states that if $\mathcal{G}$ is an expander graph, then for any small enough right vertex set V' we can always find a *closure* $\gamma\big(V'\big) \supseteq V'$ with a small kernel such that the subgraph $\mathcal{G} \setminus \gamma(V')$ has good boundary expansion. The proof of this lemma (albeit with slightly different parameters) can be found in [Raz16], but we also include it in Appendix A for completeness.

Lemma 5.3 ([Raz16]). *Let $\mathcal{G}$ be an $(s, 2)$-boundary expander. Then for every $V' \subseteq V$ with $|V'| \leq s/2$ there exists a subset $\gamma\big(V'\big) \subseteq V$ with $\gamma(V') \supseteq V'$ such that $\big|\mathrm{Ker}\big(\gamma\big(V'\big)\big)\big| \leq \big|V'\big|$ and the induced subgraph $\mathcal{G} \setminus \gamma(V')$ is an $(s/2, 1)$-boundary expander.*

Note that Lemma 5.3 is a purely existential result. We do not know how the closure is constructed and, in particular, if we want to choose closures of minimal size, then $V_1 \subseteq V_2$ does not necessarily imply $\gamma(V_1) \subseteq \gamma(V_2)$.

In order for Lemmas 5.2 and 5.3 to be useful, we need to know that there exist good enough boundary expanders. To prove this, one can just fix a left vertex set U of size m and a right vertex set V of size n and then for every $u \in U$ choose Δ neighbours from V uniformly and independently at random. A standard probabilistic argument shows that with high probability this random graph is an $m \times n$ $(s, \Delta, 2)$-boundary expander for appropriately chosen parameters. We state this formally as a lemma below. A similar lemma is proven in [Raz16] but we also provide a proof in Appendix A for the convenience of the reader.

Lemma 5.4. *There is an absolute constant $\Delta_0 \in \mathbb{N}^+$ such that for all integers Δ, s, and m satisfying $\Delta \geq \Delta_0$ and $(s\Delta)^{2\Delta} \leq m$ there exist $m \times \lceil m^{3/\Delta} \rceil$ $(s, \Delta, 2)$-boundary expanders.*

For readers familiar with expander graphs from other contexts, it might be worth pointing out that the parameters above are different from what tends to be the standard expander graph settings of $s = \Omega(m)$ and $\Delta = \mathrm{O}(1)$. Instead, in Lemma 5.4 we have s growing sublinearly in m and Δ need not be constant (although we still need $\Delta \lessapprox \log m / \log\log m$ in order to satisfy the conditions of the lemma).

In what follows, unless otherwise stated $\mathcal{G} = (U \dot\cup V, E)$ will be an $(s, 2)$-boundary expander for $s = 2k$. We will use such expanders when we do XOR substitution in our formulas as described formally in the next definition. In words, variables in the XOR formula are identified with left vertices U in $\mathcal{G}$, the pool of new variables is the right vertex set V, and every variable $u \in U$ in an XOR clause is replaced by an exclusive or $\bigoplus_{v \in N(u)} v$ over its neighbours $v \in N(u)$. We emphasize that in "standard" XOR-ification as found in the proof complexity literature all new substituted variables would be distinct, i.e., $N(u_1) \cap N(u_2) = \emptyset$ for $u_1 \neq u_2$. While this often makes formulas harder, it also increases the number of variables. Here, we use the approach in [Raz16] to instead recycle variables from a much smaller set V in the substitutions, thus decreasing the total number of variables.

Definition 5.5 (XOR substitution with recycling). Let F be an XOR formula with $Vars(F) = U$ and let $\mathcal{G} = (U \dot\cup V, E)$ be a bipartite graph. For every clause $C = (u_1, \dots, u_t, a)$ in F we let $C[\mathcal{G}]$ be the clause $(v_1^1, \dots, v_1^{z_1}, \dots, v_t^1, \dots, v_t^{z_t}, a)$, where $N(u_i) = \{v_i^1, \dots, v_i^{z_i}\}$ for all $1 \leq i \leq t$. Taking unions, we let $F[\mathcal{G}]$ be the XOR formula $F[\mathcal{G}] = \{C[\mathcal{G}] \mid C \in F\}$.

When using an $m \times m^{3/\Delta}$ $(s, \Delta, 2)$-boundary expander as in Lemma 5.4 for substitution in an m-variable XOR formula F as described in Definition 5.5, we obtain a new XOR formula $F[\mathcal{G}]$ where the number of variables have decreased significantly to $m^{3/\Delta}$. The next lemma, which is at the heart of our logic-flavoured version of hardness condensation, states that a round lower bound for the k-pebble game on F implies a round lower bound for the k-pebble game on $F[\mathcal{G}]$.

Lemma 5.6. *Let k be a positive integer and let $\mathcal{G}$ be an $m \times n$ $(2k, 2)$-boundary expander. Then if F is an XOR formula over m variables such that Player 2 wins the r-round k-pebble game on F, she also wins the $r/(2k)$-round k-pebble game on $F[\mathcal{G}]$.*

By way of comparison with [Raz16], we remark that a straightforward translation of Razborov's technique would start with formulas on which Player 1 can win with few pebbles, but needs an almost linear number of rounds to win the game, even if he has an infinite amount of pebbles.[4] Applying this without modification to Immerman's construction, we would obtain very weak bounds (and, in particular, nothing interesting for constant k). Instead, as input to our hardness condensation lemma we use a construction that has a round lower bound of $n^{1/\log k}$, and show that for hardness condensation it is not necessary that the original formula is hard over the full range.

Before embarking on a formal proof of Lemma 5.6, which is rather technical and will take the rest of this section, let us discuss the intuition behind it. The main idea to obtain a good strategy for Player 2 on the substituted formula $F[\mathcal{G}]$ is to think of the game as being played on F and simulate the survival strategy there for as long as possible (which is where boundary expansion comes into play).

Let $\mathcal{G} = (U \dot\cup V, E)$ be an $(2k, 2)$-boundary expander as stated in the lemma. We have $Vars(F) = U$ and $Vars(F[\mathcal{G}]) = V$. Given a strategy for Player 2 in the r-round k-pebble game on F, we want to convert this into a winning strategy for Player 2 for the $r/(2k)$-round k-pebble game on $F[\mathcal{G}]$. A first approach (which will not quite work) is the following.

While playing on the substituted formula $F[\mathcal{G}]$, Player 2 simulates the game on F. For every position β in the game on $F[\mathcal{G}]$, she maintains a corresponding position α on F, which is defined on all variables whose entire neighbourhood in the expander is contained in the domain of β, i.e., $Vars(\alpha) = \mathrm{Ker}(Vars(\beta))$. The assignments of α should be defined in such a way that they are *consistent with* β, i.e., so that $\alpha(u) = \bigoplus_{v \in N(u)} \beta(v)$. It then follows from the description of XORification in Definition 5.5 that α falsifies an XOR clause of F if and only if β falsifies an XOR clause of $F[\mathcal{G}]$.

Now Player 2 wants to play in such a way that if β changes to β' in one round of the game on $F[\mathcal{G}]$, then the corresponding position α also changes to α' in one round of the game on F. Intuitively, this should be done as follows. Suppose that starting from a position β, Player 1 asks for a variable $v \in V$. If v is not the last unassigned vertex in a neighbourhood of some $u \in U$, i.e., $\mathrm{Ker}(Vars(\beta)) = \mathrm{Ker}(Vars(\beta) \cup \{v\})$, then Player 2 can make an arbitrary choice as $\alpha = \alpha'$ is consistent with both choices. If v is the last free vertex in the neighbourhood of exactly one vertex u, i.e., $\{u\} = \mathrm{Ker}(Vars(\beta) \cup \{v\}) \setminus \mathrm{Ker}(Vars(\beta))$, then Player 2 assumes that she was asked for u in the simulated game on F. If in her strategy for the r-round k-pebble game on F she would answer with an assignment $a \in \{0, 1\}$ which would yield the new position $\alpha' = \alpha \cup \{u \mapsto a\}$, then in the game on $F[\mathcal{G}]$ she now sets v to the right value $b \in \{0, 1\}$ so that the new position $\beta' = \beta \cup \{v \mapsto b\}$ satisfies the consistency property $\alpha'(u) = \bigoplus_{v \in N(u)} \beta'(v)$. If Player 2 could follow this strategy, then the number of rounds she would survive the game on $F[\mathcal{G}]$ would be lower-bounded by the number of rounds she survives in the game on F.

There is a gap in this intuitive argument, however, namely how to handle the case when the queried variable v completes the neighbourhood of two (or more) vertices u_1, u_2 at the same time. If it holds that

[4] In terms of resolution, this corresponds to formulas that are refutable in small width, but where every resolution refutation has almost linear depth.

$\{u_1, u_2\} \subseteq \mathrm{Ker}(\mathit{Vars}(\beta) \cup \{v\}) \setminus \mathrm{Ker}(\mathit{Vars}(\beta))$, then we have serious problems. Following the strategy above for u_1 and u_2 separately can yield two different and conflicting ways of assigning v, meaning that for the new position β' there will be no consistent assignment α' of $\mathrm{Ker}(\mathit{Vars}(\beta'))$.

To circumvent this problem and implement the proof idea above, we will use the boundary expansion of $\mathcal{G}$ to ensure that this problematic case does not occur. For instance, suppose that the graph $\mathcal{G}' = \mathcal{G} \setminus \mathit{Vars}(\beta)$, which is the induced subgraph of $\mathcal{G}$ on $U \setminus \mathit{Vars}(\alpha)$ and $V \setminus \mathit{Vars}(\beta)$, has boundary expansion at least 1. Then the bad situation described above with two variables u_1, u_2 having neighbourhood $N^{\mathcal{G}'}(u_1) = N^{\mathcal{G}'}(u_2) = \{v\}$ in $\mathcal{G}'$ cannot arise, since this would imply $\partial^{\mathcal{G}'}(\{u_1, u_2\}) = \emptyset$, contradicting the expansion properties of $\mathcal{G}'$. Unfortunately, we cannot ensure boundary expansion of $\mathcal{G} \setminus \mathit{Vars}(\beta)$ for every position β, but we can apply Lemma 5.3 and extend the current position to a larger one that is defined on $\gamma(\mathit{Vars}(\beta))$ and has the desired expansion property. Since Lemma 5.3 ensures that the domain $\mathrm{Ker}(\gamma(\mathit{Vars}(\beta)))$ of our assignment α under construction is bounded by $|\alpha| \leq |\beta| \leq k$, such an extension will still be good enough.

We now proceed to present a formal proof. When doing so, it turns out to be convenient for us to prove the contrapositive of the statement discussed above. That is, instead of transforming a strategy for Player 2 in the r-round k-pebble game on F to a strategy for the $r/(2k)$-round k-pebble game on $F[\mathcal{G}]$ for an $(2k, 2)$-boundary expander $\mathcal{G}$, we will show that a winning strategy for Player 1 in the game on the substituted formula $F[\mathcal{G}]$ can be used to obtain a winning strategy for Player 1 in the game on the original formula F.

Suppose that β is any position in the k-pebble game on $F[\mathcal{G}]$, i.e., a partial assignment of variables in V. Since $\mathcal{G}$ is a $(2k, 2)$-boundary expander and $|\beta| \leq k$, we can apply Lemma 5.3 to obtain a superset $\gamma(\mathit{Vars}(\beta)) \supseteq \mathit{Vars}(\beta)$ having the properties that $|\mathrm{Ker}(\gamma(\mathit{Vars}(\beta)))| \leq |\mathit{Vars}(\beta)|$ and the induced subgraph $\mathcal{G} \setminus \gamma(\mathit{Vars}(\beta))$ is a $(k, 1)$-boundary expander. For the rest of this section, fix a minimal such set $\gamma(V')$ for every $V' = \mathit{Vars}(\beta)$ corresponding to a position β in the k-pebble game. This will allow us to define formally what we mean by *consistent* positions in the two games on F and $F[\mathcal{G}]$ as described next.

Definition 5.7 (Consistent positions). Let α be a position in the pebble game on F, i.e., a partial assignment of variables in U, and let β be a partial assignment of variables in V corresponding to a position in the pebble game on $F[\mathcal{G}]$. We say that α *is consistent with* β if there exists an extension $\beta_{\text{ext}} \supseteq \beta$ with $\mathit{Vars}(\beta_{\text{ext}}) = N\big(\mathit{Vars}(\alpha)\big) \cup \mathit{Vars}(\beta)$ such that for all $u \in \mathit{Vars}(\alpha)$ it holds that $\alpha(u) = \bigoplus_{v \in N(u)} \beta_{\text{ext}}(v)$.

Let β be a position in the k-pebble game on the XOR-substituted formula $F[\mathcal{G}]$ and let $\gamma(V')$ be the fixed, minimal closure of β chosen above. Then we let $\mathit{Cons}(\beta)$ denote the set of all positions α with $\mathit{Vars}(\alpha) = \mathrm{Ker}(\gamma(\mathit{Vars}(\beta)))$ in the pebble game on F that are consistent with β.

Observe that for $\alpha_1 \subseteq \alpha$ and $\beta_1 \subseteq \beta_2$ it holds that if α_2 is consistent with β_1 then so is α_1, and if α_1 is consistent with β_2 then α_1 is consistent also with β_1. Furthermore, by Lemma 5.3 we have $|\alpha| \leq |\beta|$ for all $\alpha \in \mathit{Cons}(\beta)$. The next claim states the core inductive argument.

Claim 5.8. Let β be a position on $F[\mathcal{G}]$ for an $(2k, 2)$-boundary expander $\mathcal{G}$ and suppose that Player 1 wins the i-round k-pebble game on $F[\mathcal{G}]$ from position β. Then Player 1 has a strategy to win the k-pebble game on F within $2ki$ rounds from every position $\alpha \in \mathit{Cons}(\beta)$.

We note that this claim is just a stronger version of (the contrapositive of) Lemma 5.6.

Proof of Lemma 5.6 assuming Claim 5.8. Note that if $r/(2k) < 1$, then the lemma is trivially true, as Player 1 always needs at least one round to win the pebble game from the empty position. Otherwise, we apply Claim 5.8 with parameters $\beta = \emptyset$ and $i = r/(2k)$. Since $\mathit{Cons}(\emptyset) = \{\emptyset\}$, we directly get the contrapositive statement of Lemma 5.6 that if Player 1 wins the $r/(2k)$-round k-pebble game on $F[\mathcal{G}]$, then he wins the r-round k-pebble game on F. □

All that remains for us to do now is to establish Claim 5.8, after which the hardness condensation lemma will follow easily.

Proof of Claim 5.8. The proof is by induction on i. For the base case $i = 0$ we have to show that if β falsifies an XOR clause in $F[\mathcal{G}]$, then every assignment $\alpha \in \mathit{Cons}(\beta)$ falsifies an XOR clause in F. But if β falsifies a clause of $F[\mathcal{G}]$, which by construction has the form $C[\mathcal{G}]$ for some clause C from F, then by Definitions 5.5 and 5.7 it holds that every $\alpha \in \mathit{Cons}(\beta)$ falsifies C.

For the induction step, suppose that the statement holds for $i-1$ and assume that Player 1 wins the i-round k-pebble game on $F[\mathcal{G}]$ from position β. The ith round consists of two steps:

1. Player 1 first chooses a subassignment $\beta' \subseteq \beta$.
2. He then asks for the value of one variable $v \in V \setminus \mathit{Vars}(\beta')$, to which Player 2 chooses an assignment $b \in \{0,1\}$ yielding the new position $\beta' \cup \{v \mapsto b\}$.

As Player 1 has a strategy to win from β within i rounds, it follows that he can win from both $\beta' \cup \{v \mapsto 0\}$ and $\beta' \cup \{v \mapsto 1\}$ within $i-1$ rounds. By the inductive assumption we then deduce for the set of assignments

$$\mathit{Cons}(\beta' * v) := \mathit{Cons}(\beta' \cup \{v \mapsto 0\}) \cup \mathit{Cons}(\beta' \cup \{v \mapsto 1\}) \tag{5.2}$$

consistent with either $\beta' \cup \{v \mapsto 0\}$ or $\beta' \cup \{v \mapsto 1\}$ that the following statement holds.

Subclaim 5.9. Player 1 can win the k-pebble game on F within $2k(i-1)$ rounds from all positions in $\mathit{Cons}(\beta' * v)$.

Note that a position is in $\mathit{Cons}(\beta' * v)$ if it is consistent with either $\beta' \cup \{v \mapsto 0\}$ or $\beta' \cup \{v \mapsto 1\}$. Therefore, $\mathit{Cons}(\beta' * v)$ is the set of all positions over $\mathrm{Ker}(\gamma(\beta') \cup \{v\})$ that are consistent with β'. What remains to show is that from every position $\alpha \in \mathit{Cons}(\beta)$ Player 1 can reach some position in $\mathit{Cons}(\beta' * v)$ within $2k$ rounds. We split the proof into two steps, corresponding to the two steps in the move of Player 1 from position β.

Subclaim 5.10. From every position $\alpha \in \mathit{Cons}(\beta)$ Player 1 can reach some position in $\mathit{Cons}(\beta')$ for $\beta' \subseteq \beta$ within k rounds.

Subclaim 5.11. From every position $\alpha \in \mathit{Cons}(\beta')$ Player 1 can reach some position in $\mathit{Cons}(\beta' * v)$ within k rounds.

Let us establish Subclaims 5.10 and 5.11 in reverse order.

Proof of Subclaim 5.11. Player 1 starts with an assignment $\alpha_{\text{start}} \in \mathit{Cons}(\beta')$, which is defined over the variables $U_{\text{start}} = \mathrm{Ker}(\gamma(\mathit{Vars}(\beta')))$, and wants to reach some assignment $\alpha_{\text{end}} \in \mathit{Cons}(\beta' * v)$ defined over the variables $U_{\text{end}} = \mathrm{Ker}(\gamma(\mathit{Vars}(\beta') \cup \{v\}))$.

If $\mathrm{Ker}(\gamma(\mathit{Vars}(\beta'))) = \mathrm{Ker}(\gamma(\mathit{Vars}(\beta') \cup \{v\}))$, then Player 1 can choose $\alpha_{\text{end}} = \alpha_{\text{start}}$. To see this, note that if α_{start} assigns a value to some $u \in N(v)$, then since $\alpha_{\text{start}} \in \mathit{Cons}(\beta')$ it holds by Definition 5.7 that $N(u) \subseteq \gamma(\mathit{Vars}(\beta'))$, and thus α_{start} is already consistent with $\beta' \cup \{v \mapsto b\}$ for some $b \in \{0,1\}$. Hence, Player 1 need not ask any question in this case, but the induction hypothesis immediately yields the desired conclusion.

The more interesting case is when $\mathrm{Ker}(\gamma(\mathit{Vars}(\beta'))) \neq \mathrm{Ker}(\gamma(\mathit{Vars}(\beta') \cup \{v\}))$. Now Player 1 first deletes all assignments of variables in $U_{\text{start}} \setminus U_{\text{end}}$ from α_{start} to get α_0. Since $\alpha_0 \subseteq \alpha_{\text{start}}$ and α_{start} is consistent with β' by assumption, α_0 is also consistent with β'. Afterwards, he asks for all variables in $U' = U_{\text{end}} \setminus U_{\text{start}}$. We need to argue that regardless of how Player 2 answers, it holds that Player 1 reaches a position that is consistent with β' .This is where the peeling argument in Lemma 5.2 is needed.

As discussed above, by our choice of the closure $\gamma(\mathit{Vars}(\beta'))$ (obtained using Lemma 5.3) we know that the bipartite graph $\mathcal{G}' = \mathcal{G} \setminus \gamma(\mathit{Vars}(\beta'))$ is a $(k,1)$-boundary expander and furthermore that for $U' = U_{\text{end}} \setminus U_{\text{start}}$ it holds that $|U'| \leq |U_{\text{end}}| \leq |\mathit{Vars}(\beta') \cup \{v\}| \leq k$, as observed right after Definition 5.7. Hence, we can apply Lemma 5.2 to $\mathcal{G}'$ and U' to get an ordered sequence $u_1, \ldots, u_t$ satisfying $N^{\mathcal{G}'}(u_i) \setminus N^{\mathcal{G}'}(\{u_1, \ldots, u_{i-1}\}) \neq \emptyset$. We will think of Player 1 as querying the (at most k) vertices in U' in this order, after which he ends up with a position α_{end} defined on the variables U_{end}.

To argue that the position α_{end} obtained in this way is consistent with β' independently of how Player 2 answers, and is hence contained in $\mathit{Cons}\big(\beta' * v\big)$, we show inductively that all positions encountered during the transition from α_{start} to α_{end} are consistent with β'. As already noted, this holds for the position α_0 obtained from α_{start} by deleting all assignments of variables in $U_{\text{start}} \setminus U_{\text{end}}$. For the induction step, let $i \geq 0$ and assume inductively that the current position α_i over

$$U_i := (U_{\text{start}} \cap U_{\text{end}}) \cup \{u_j \mid 1 \leq j \leq i\} \tag{5.3}$$

is consistent with β'. Now Player 1 asks about the variable u_{i+1} and Player 2 answers with a value a_{i+1}. Since α_i is consistent with β', there is an assignment $\beta_{\text{ext}} \supseteq \beta'$ that sets the variables $v \in N(\mathit{Vars}(\alpha_i))$ to the right values such that $\alpha_i(u) = \bigoplus_{v \in N(u)} \beta_{\text{ext}}(v)$ for all $u \in \mathit{Vars}(\alpha_i)$. By our ordering of $U' = \{u_1, \ldots, u_t\}$ chosen above we know that u_{i+1} has at least one neighbour on the right-hand side V that is neither contained in $N^{\mathcal{G}}(U_i) = N^{\mathcal{G}}(\mathit{Vars}(\alpha_i))$ nor in the domain of β'. Hence, regardless of which value a_{i+1} Player 2 chooses for her answer we can extend the assignment β_{ext} to the variables $N^{\mathcal{G}}(u_{i+1}) \setminus \big(N^{\mathcal{G}}\big(\mathit{Vars}(\alpha_i)\big) \cup \mathit{Vars}\big(\beta'\big)\big)$ in such a way that $\bigoplus_{v \in N(u_{i+1})} \beta_{\text{ext}}(v) = a_{i+1}$. This shows that α_{i+1} defined over $U_{i+1} = (U_{\text{start}} \cap U_{\text{end}}) \cup \{u_j \mid 1 \leq j \leq i+1\}$ is consistent with β'. Subclaim 5.11 now follows by the induction principle. ⊣

Before proving Subclaim 5.10, we should perhaps point out why this claim is not vacuous. Recalling the discussion just below Lemma 5.3, this is because the condition $V_1 \subseteq V_2$ does not allow us to conclude that $\gamma(V_1) \subseteq \gamma(V_2)$.

Proof of Subclaim 5.10. The proof is similar to that of Subclaim 5.11 above. Player 1 starts with an assignment $\alpha_{\text{start}} \in \mathit{Cons}(\beta)$ and wants to reach some assignment in $\mathit{Cons}(\beta')$ for $\beta' \subseteq \beta$ within k rounds. By assumption, α_{start} is consistent with β and therefore (since $\beta' \subseteq \beta$) is also consistent with β'. Player 1 deletes all assignments from the domain $U_{\text{start}} = \mathrm{Ker}(\gamma(\mathit{Vars}(\beta)))$ of α_{start} that do not occur in the domain $U_{\text{end}} = \mathrm{Ker}(\gamma(\mathit{Vars}(\beta')))$ of positions in $\mathit{Cons}(\beta')$, resulting in the position $\alpha_0 \subseteq \alpha_{\text{start}}$ that is consistent with β'. Next, he applies Lemma 5.2 to $\mathcal{G}' = \mathcal{G} \setminus \gamma(\mathit{Vars}(\beta))$ to obtain an ordering of the remaining variables $U_{\text{end}} \setminus U_{\text{start}}$. In the same way as above he can query the variables in this order while maintaining the invariant that the current position is consistent with β'. ⊣

Combining Subclaims 5.9, 5.10 and 5.11, we conclude that Player 1 wins from every position $\alpha \in \mathit{Cons}(\beta)$ within $2ki$ rounds. This concludes the proof of Claim 5.8. □

We are finally in a position to give a formal proof of Lemma 3.3.

Proof of Lemma 3.3. Let $\Delta_0 \in \mathbb{N}^+$ be the constant in Lemma 5.4. Suppose we are given an m-variable p-XOR formula F and parameters ℓ_{lo}, ℓ_{hi}, r, Δ satisfying the conditions in Lemma 3.3 that $\ell_{\text{hi}}/\ell_{\text{lo}} \geq \Delta \geq \Delta_0$ and $(2\ell_{\text{hi}}\Delta)^{2\Delta} \leq m$.

Fix $k := \ell_{\text{hi}}$ and $s := 2\ell_{\text{hi}}$. Since $(s\Delta)^{2\Delta} \leq m$ and $\Delta \geq \Delta_0$, we appeal to Lemma 5.4 to obtain an $m \times \lceil m^{3/\Delta} \rceil$ $(s, \Delta, 2)$-boundary expander $\mathcal{G} = (U \dot\cup V, E)$, and applying XORification with respect to $\mathcal{G}$ we construct the formula $H := F[\mathcal{G}]$. Clearly, H is an (Δp)-XOR formula with $\lceil m^{3/\Delta} \rceil$ variables. We want to prove that Player 1 has a winning strategy for the $(\Delta\ell_{\text{lo}})$-pebble game on H as guaranteed by Lemma 3.3(a'), but that he does not win the ℓ_{hi}-pebble game on H within $r/(2\ell_{\text{hi}})$ rounds as stated in Lemma 3.3(b').

For the upper bound in Lemma 3.3(a'), we recall that Player 1 has a winning strategy in the ℓ_{lo}-pebble game on F by assumption (a) in the lemma. He can use this strategy to win the $(\Delta\ell_{\text{lo}})$-pebble game on H as follows. Whenever his strategy tells him to ask for a variable $u \in U = \mathit{Vars}(F)$, he instead asks for the at most Δ variables in $N(u) \subseteq V = \mathit{Vars}(H)$ and assigns to u the value that corresponds to the parity of the answers Player 2 gives for $N(u)$. In this way, he can simulate his strategy on F until he reaches an assignment that contradicts an XOR clause C from F. As the corresponding assignment of the variables $\{v \mid v \in N(u), u \in \mathit{Vars}(C)\}$ falsifies the constraint $C[\mathcal{G}] \in H$, at this point Player 1 wins the $(\Delta\ell_{\text{lo}})$-pebble game on H.

The lower bound in Lemma 3.3(b') follows immediately from Lemma 5.6. By assumption (b) in Lemma 3.3, Player 1 does not win the ℓ_{hi}-pebble game on F within r rounds. Since $\mathcal{G}$ is an $m \times n$ $(2k, 2)$-boundary expander, Lemma 5.6 says that that he does not win the ℓ_{hi}-pebble game on $H = F[\mathcal{G}]$ within $r/(2\ell_{\text{hi}})$ rounds either. This concludes the proof of Lemma 3.3 □

6 Concluding Remarks

In this paper we prove an $n^{\Omega(k/\log k)}$ lower bound on the minimal quantifier depth of L^k and C^k sentences that distinguish two finite n-element relational structures, nearly matching the trivial n^{k-1} upper bound. By the known connection to the k-dimensional Weisfeiler–Leman algorithm, this result implies near-optimal $n^{\Omega(k/\log k)}$ lower bounds also on the number of refinement steps of this algorithm. The key technical ingredient in our proof is the hardness condensation technique recently introduced by Razborov [Raz16] in the context of proof complexity, which we translate into the language of finite variable logics and use to reduce the domain size of relational structures while maintaining the minimal quantifier depth required to distinguish them.

An obvious open problem is to improve our lower bound. One way to achieve this would be to strengthen the lower bound on the number of rounds in the k-pebble game on 3-XOR formulas in Lemma 3.2 from $n^{1/\log k}$ to n^{δ} for some $\delta \gg 1/\log k$. By the hardness condensation lemma this would directly improve our lower bound from $n^{\Omega(k/\log k)}$ to $n^{\Omega(\delta k)}$.

The structures on which our lower bounds hold are n-element relational structures of arity $\Theta(k)$ and size $n^{\Theta(k)}$. We would have liked to have this results also for structures of bounded arity, such as graphs. However, the increase of the arity is inherent in the method of amplifying hardness by making XOR substitutions. An optimal lower bound of $n^{\Omega(k)}$ on the quantifier depth required to distinguish two n-vertex graphs has been obtained by the first author in an earlier work [Ber14] for the *existential-positive fragment* of L^k. Determining the quantifier rank of full L^k and C^k on n-vertex graphs remains an open problem.

Another open question related to our results concerns the complexity of finite variable equivalence for non-constant k. What is the complexity of deciding, given two structures and a parameter k, whether the structures are equivalent in L^k or C^k? As this problem can be solved in time $(\|\mathcal{A}\| + \|\mathcal{B}\|)^{O(k)}$, it is in EXPTIME if k is part of the input. It has been conjectured that this problem is EXPTIME-complete [GKL+07], but it is not even known whether it is NP-hard. Note that the quantifier depth is connected to the computational complexity of the equivalence problem by the fact that an upper bound of the form $n^{\mathrm{O}(1)}$ on n-element structures would have implied that testing equivalence is in PSPACE. Hence, our lower bounds on the quantifier depth can be seen as a necessary requirement for establishing EXPTIME-hardness of the equivalence problem.

Acknowledgements

We are very grateful to Alexander Razborov for patiently explaining the hardness condensation technique in [Raz16] during numerous and detailed discussions. We also want to thank Neil Immerman for helping us find relevant references to previous works and explaining some of the finer details in [Imm81]. The first author wishes to acknowledge useful feedback from the participants of Dagstuhl Seminar 15511 *The Graph Isomorphism Problem*. Finally, we are most indebted to the anonymous *LICS* reviewers for very detailed comments that helped improve the manuscript considerably.

Part of the work of the first author was performed while at KTH Royal Institute of Technology supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD). The research of the second author was supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007–2013) / ERC grant agreement no. 279611 and by Swedish Research Council grants 621-2010-4797 and 621-2012-5645.

References

[AD08] Albert Atserias and Víctor Dalmau. A combinatorial characterization of resolution width. *Journal of Computer and System Sciences*, 74(3):323–334, May 2008. Preliminary version in *CCC '03*.

[AR03] Michael Alekhnovich and Alexander A. Razborov. Lower bounds for polynomial calculus: Non-binomial case. *Proceedings of the Steklov Institute of Mathematics*, 242:18–35, 2003. Available at http://people.cs.uchicago.edu/~razborov/files/misha.pdf. Preliminary version in *FOCS '01*.

[Bab16] László Babai. Graph isomorphism in quasipolynomial time. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC '16)*, pages 684–697, June 2016.

[Bar77] Jon Barwise. On Moschovakis closure ordinals. *Journal of Symbolic Logic*, 42(2):292–296, June 1977.

[Ben02a] Eli Ben-Sasson. Hard examples for the bounded depth Frege proof system. *Computational Complexity*, 11(3-4):109–136, 2002.

[Ben02b] Eli Ben-Sasson. Size space tradeoffs for resolution. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC '02)*, pages 457–464, May 2002. Journal version in [Ben09].

[Ben09] Eli Ben-Sasson. Size-space tradeoffs for resolution. *SIAM Journal on Computing*, 38(6):2511–2525, May 2009. Preliminary version in *STOC '02*.

[Ber14] Christoph Berkholz. The propagation depth of local consistency. In *Proceedings of the 20th International Conference on Principles and Practice of Constraint Programming (CP '14)*, pages 158–173, September 2014.

[BGIP01] Samuel R. Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. Linear gaps between degrees for the polynomial calculus modulo distinct primes. *Journal of Computer and System Sciences*, 62(2):267–289, March 2001. Preliminary version in *CCC '99*.

[BN08] Eli Ben-Sasson and Jakob Nordström. Short proofs may be spacious: An optimal separation of space and length in resolution. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '08)*, pages 709–718, October 2008.

[BN11] Eli Ben-Sasson and Jakob Nordström. Understanding space in proof complexity: Separations and trade-offs via substitutions. In *Proceedings of the 2nd Symposium on Innovations in Computer Science (ICS '11)*, pages 401–416, January 2011.

[BN16] Christoph Berkholz and Jakob Nordström. Supercritical space-width trade-offs for resolution. In *Proceedings of the 43rd International Colloquium on Automata, Languages and Programming (ICALP '16)*, volume 55 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 57:1–57:14, July 2016.

[BNT13] Chris Beck, Jakob Nordström, and Bangsheng Tang. Some trade-off results for polynomial calculus. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing (STOC '13)*, pages 813–822, May 2013.

[BP07] Joshua Buresh-Oppenheim and Toniann Pitassi. The complexity of resolution refinements. *Journal of Symbolic Logic*, 72(4):1336–1352, December 2007. Preliminary version in *LICS '03*.

[BW01] Eli Ben-Sasson and Avi Wigderson. Short proofs are narrow—resolution made simple. *Journal of the ACM*, 48(2):149–169, March 2001. Preliminary version in *STOC '99*.

[CFI92] Jin-Yi Cai, Martin Fürer, and Neil Immerman. An optimal lower bound on the number of variables for graph identifications. *Combinatorica*, 12(4):389–410, December 1992.

[Coo74] Stephen A. Cook. An observation on time-storage trade off. *Journal of Computer and System Sciences*, 9(3):308–316, 1974. Preliminary version in *STOC '73*.

[FLM+13] Yuval Filmus, Massimo Lauria, Mladen Mikša, Jakob Nordström, and Marc Vinyals. Towards an understanding of polynomial calculus: New separations and lower bounds (Extended abstract). In *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP '13)*, volume 7965 of *Lecture Notes in Computer Science*, pages 437–448. Springer, July 2013.

[Fü01] Martin Fürer. Weisfeiler–Lehman refinement requires at least a linear number of iterations. In *Proceedings of the 28th International Colloquium on Automata, Languages, and Programming (ICALP '01)*, volume 2076 of *Lecture Notes in Computer Science*, pages 322–333. Springer, July 2001.

[GKL+07] Erich Grädel, Phokion G. Kolaitis, Leonid Libkin, Maarten Marx, Joel Spencer, Moshe Y. Vardi, Yde Venema, and Scott Weinstein. *Finite Model Theory and Its Applications*. Springer, 2007.

[Gri01] Dima Grigoriev. Linear lower bound on degrees of Positivstellensatz calculus proofs for the parity. *Theoretical Computer Science*, 259(1–2):613–622, May 2001.

[Gro98] Martin Grohe. Finite variable logics in descriptive complexity theory. *Bulletin of Symbolic Logic*, 4(4):345–398, 1998.

[Gro99] Martin Grohe. Equivalence in finite-variable logics is complete for polynomial time. *Combinatorica*, 19(4):507–532, October 1999.

[Gro12] Martin Grohe. Fixed-point definability and polynomial time on graphs with excluded minors. *Journal of the ACM*, 59(5):27:1–27:64, October 2012. Preliminary version in *LICS '10*.

[GS96] Yuri Gurevich and Saharon Shelah. On finite rigid structures. *Journal of Symbolic Logic*, 61(2):549–562, June 1996.

[Hel96] Lauri Hella. Logical hierarchies in PTIME. *Information and Computation*, 129:1–19, August 1996.

[IL90] Neil Immerman and Eric Lander. Describing graphs: a first-order approach to graph canonization. In Alan L. Selman, editor, *Complexity Theory Retrospective: In Honor of Juris Hartmanis on the Occasion of His Sixtieth Birthday*, pages 59–81. Springer, 1990.

[Imm81] Neil Immerman. Number of quantifiers is better than number of tape cells. *Journal of Computer and System Sciences*, 22(3):384–406, June 1981.

[Imm82] Neil Immerman. Upper and lower bounds for first order expressibility. *Journal of Computer and System Sciences*, 25(1):76–98, August 1982.

[KI06] Arist Kojevnikov and Dmitry Itsykson. Lower bounds of static Lovász–Schrijver calculus proofs for Tseitin tautologies. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP '06)*, volume 4051 of *Lecture Notes in Computer Science*, pages 323–334. Springer, July 2006.

[KS16] Sandra Kiefer and Pascal Schweitzer. Upper bounds on the quantifier depth for graph differentiation in first order logic. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '16)*, pages 287–296, July 2016.

[KV15] Andreas Krebs and Oleg Verbitsky. Universal covers, color refinement, and two-variable counting logic: Lower bounds for the depth. In *Proceedings of the 30th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '15)*, pages 689–700, July 2015.

[MN15] Mladen Mikša and Jakob Nordström. A generalized method for proving polynomial calculus degree lower bounds. In *Proceedings of the 30th Annual Computational Complexity Conference (CCC '15)*, volume 33 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 467–487, June 2015.

[Nor13] Jakob Nordström. Pebble games, proof complexity and time-space trade-offs. *Logical Methods in Computer Science*, 9:15:1–15:63, September 2013.

[Raz16] Alexander A. Razborov. A new kind of tradeoffs in propositional proof complexity. *Journal of the ACM*, 63:16:1–16:14, April 2016.

[Sch08] Grant Schoenebeck. Linear level Lasserre lower bounds for certain k-CSPs. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS '08)*, pages 593–602, October 2008.

[Urq87] Alasdair Urquhart. Hard examples for resolution. *Journal of the ACM*, 34(1):209–219, January 1987.

[Var95] Moshe Y. Vardi. On the complexity of bounded-variable queries (Extended abstract). In *Proceedings of the 14th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS '95)*, pages 266–276, May 1995.

A Existence and Properties of Expander Graphs

In this appendix we present proofs of Lemmas 5.3 and 5.4, starting with the latter lemma. We again remark that most of this material can already be found in a similar form in [Raz16], although the exact parameters are somewhat different. It also seems appropriate to point out that there is a significant overlap with essentially identical technical lemmas in [BN16].

Just to avoid ambiguity, let us state explicitly that even though we have the Euler number e appearing below, we still think of all logarithms as being taken to base 2 (though this should not really matter).

Lemma 5.4 (restated). *There is an absolute constant $\Delta_0 \in \mathbb{N}^+$ such that for all Δ, s, m satisfying $\Delta \geq \Delta_0$ and $(s\Delta)^{2\Delta} \leq m$ there exist $m \times \lceil m^{3/\Delta} \rceil$ $(s, \Delta, 2)$-boundary expanders.*

Proof. Let U and V be two disjoint sets of vertices of size $|U| = m$ and $|V| = n = \lceil m^{3/\Delta} \rceil$. For every $u \in U$ we choose Δ times a neighbour $v \in V$ uniformly at random with repetitions. This yields a bipartite graph $\mathcal{G} = (U \dot\cup V, E)$ of left-degree at most Δ. In the sequel we show that $\mathcal{G}$ is likely to be an $(s, \Delta, 2)$-boundary expander.

First note that for every set $U' \subseteq U$ all neighbours $v \in \mathcal{N}(U') \setminus \partial(U')$ that are not in the boundary of U' have at least two neighbours in U'. Since there are at most $\Delta|U'| - |\partial(U')|$ edges between U' and $\mathcal{N}(U') \setminus \partial(U')$, it follows that $|\mathcal{N}(U') \setminus \partial(U')| \leq (\Delta|U'| - |\partial(U')|)/2$ and hence

$$|\mathcal{N}(U')| \leq \frac{|\partial(U')| + \Delta|U'|}{2} \ . \tag{A.1}$$

If $\mathcal{G}$ is not an $(s, \Delta, 2)$-boundary expander, then there is a set U' of size $\left|U'\right| = \ell \leq s$ that has a boundary $\partial(U')$ of size $\left|\partial(U')\right| < 2\ell$ and from (A.1) it then follows that $|\mathcal{N}(U')| < (1 + \Delta/2)\ell$. By a

union bound argument (and relaxing to non-strict inequalities) we obtain

$$\Pr[\mathcal{G} \text{ is not an } (s,\Delta,2)\text{-boundary expander}] \tag{A.2a}$$
$$\leq \sum_{\ell=1}^{s} \sum_{U'\subseteq U;\,|U'|=\ell} \Pr\left[|\partial(U')| \leq 2\ell\right] \tag{A.2b}$$
$$\leq \sum_{\ell=1}^{s} \sum_{U'\subseteq U;\,|U'|=\ell} \Pr\left[|\mathcal{N}(U')| \leq (1+\Delta/2)\ell\right] \tag{A.2c}$$
$$\leq \sum_{\ell=1}^{s} \binom{m}{\ell}\binom{n}{(1+\Delta/2)\ell}\left(\frac{(1+\Delta/2)\ell}{n}\right)^{\Delta\ell} \tag{A.2d}$$
$$\leq \sum_{\ell=1}^{s} m^{\ell}\left(\frac{en}{(1+\Delta/2)\ell}\right)^{(1+\Delta/2)\ell}\left(\frac{(1+\Delta/2)\ell}{n}\right)^{\Delta\ell} \tag{A.2e}$$
$$= \sum_{\ell=1}^{s} m^{\ell}(en)^{(1+\Delta/2)\ell}\left((1+\Delta/2)\ell\right)^{(\Delta/2-1)\ell} n^{-\Delta\ell} \tag{A.2f}$$
$$\leq \sum_{\ell=1}^{s} n^{(\Delta/3)\ell}(en)^{(1+\Delta/2)\ell}\left((1+\Delta/2)\ell\right)^{(\Delta/2-1)\ell} n^{-\Delta\ell} \tag{A.2g}$$
$$= \sum_{\ell=1}^{s} n^{(\Delta/3)\ell} n^{\frac{\log e}{\log n}(1+\Delta/2)\ell} n^{\frac{1}{\log n}\log\left((\Delta/2+1)\ell\right)(\Delta/2-1)\ell} n^{(-\Delta/2+1)\ell} \tag{A.2h}$$
$$\leq \sum_{\ell=1}^{s} n^{\left(\frac{\log e}{\log n}\Delta+\frac{1}{\log n}\log(\Delta s)(\Delta/2-1)-\Delta/6+1\right)\ell}\ , \tag{A.2i}$$

where in going from (A.2d) to (A.2e) we use the inequality $\binom{n}{k} \leq \left(\frac{en}{k}\right)^k$ for $e \approx 2.718$ denoting the Euler number, and in going from (A.2h) to (A.2i) we assume that $\Delta \geq 2$ and also use that $\ell \leq s$.

In order to show that the expression (A.2i) is bounded away from 1—which implies that $\mathcal{G}$ is an $(s,\Delta,2)$-boundary expander with constant probability—it suffices to study the exponent and prove that there is a constant $\varepsilon > 0$ such that

$$\frac{\log e}{\log n}\Delta + \frac{1}{\log n}\log(\Delta s)\left(\frac{\Delta}{2}-1\right) - \frac{\Delta}{6} + 1 \leq -\varepsilon < 0\ , \tag{A.3}$$

which holds if there is a constant $\varepsilon' = \varepsilon/\Delta$ such that

$$\frac{\log e}{\log n} + \frac{1}{\log n}\log(\Delta s)\left(\frac{1}{2}-\frac{1}{\Delta}\right) - \frac{1}{6} + \frac{1}{\Delta} \leq -\varepsilon' < 0\ . \tag{A.4}$$

Since $(s\Delta)^{2\Delta} \leq m \leq n^{\Delta/3}$ we have $s\Delta \leq n^{1/6}$, and it follows that

$$\frac{\log e}{\log n} + \frac{\log(\Delta s)(1/2-1/\Delta)}{\log n} - \frac{1}{6} + \frac{1}{\Delta} \tag{A.5a}$$
$$\leq \frac{\log e}{\log n} + \frac{\log\left(n^{1/6}\right)}{2\log n} - \frac{1}{6} + \frac{1}{\Delta} \tag{A.5b}$$
$$= \frac{\log e}{\log n} - \frac{1}{12} + \frac{1}{\Delta} \tag{A.5c}$$
$$\leq -\varepsilon' < 0 \tag{A.5d}$$

where we can make the last inequality hold for ε' small enough and n and Δ large enough. Fix ε' satisfying $0 < \varepsilon' < 1/12$ and choose n_0 so that the inequality between (A.5c) and (A.5d) holds for

any $n \geq n_0$ and $\Delta \geq 13$. Then we obtain that (A.2i) is bounded by $\sum_{\ell=1}^{s} n^{-\varepsilon'\ell}$. Insisting in addition that $n \geq 3^{1/\varepsilon'}$, we can upper-bound (A.2i) by

$$\sum_{\ell=1}^{s} n^{-\varepsilon'\ell} \leq \sum_{\ell=1}^{\infty} \left(\tfrac{1}{3}\right)^{\ell} \leq \tfrac{1}{2} \ . \tag{A.6}$$

It remains to calculate how to set Δ_0 to make sure that all of these conditions hold. Note that by assumption we have $(s\Delta)^{2\Delta} \leq m$, which implies that $\Delta^{\Delta} \leq m$. It follows that we will always have $n = \lceil m^{3/\Delta} \rceil \geq (\Delta^{\Delta})^{3/\Delta} = \Delta^3 \geq (\Delta_0)^3$ and hence it is sufficient to choose $\Delta_0 \geq \max\left({n_0}^{1/3}, 3^{1/3\varepsilon'}, 13\right)$. This concludes the proof of the lemma. □

We next prove that in a good enough boundary expander it holds that for any small enough right vertex set V' there is a superset $\gamma(V') \supseteq V'$ with a small kernel such that the induced subgraph $\mathcal{G} \setminus \gamma(V')$ (obtained from $\mathcal{G}$ by deleting V' and then all isolated vertices from U) is also a good boundary expander. Recall that we refer to this set $\gamma(V')$ as the *closure* of V'.

Lemma 5.3 (restated). *Let $\mathcal{G}$ be an $(s, 2)$-boundary expander. Then for every $V' \subseteq V$ with $|V'| \leq s/2$ there exists a subset $\gamma(V') \subseteq V$ with $\gamma(V') \supseteq V'$ such that $\left|\mathrm{Ker}(\gamma(V'))\right| \leq \left|V'\right|$ and the induced subgraph $\mathcal{G} \setminus \gamma(V')$ is an $(s/2, 1)$-boundary expander.*

Proof. Let $\mathcal{G} = (U \dot\cup V, E)$ be an $(s, 2)$-boundary expander and let $V' \subseteq V$ have size $|V'| \leq s/2$. We construct an increasing sequence $V' = V_0 \subset V_1 \subset \cdots \subset V_\tau = \gamma(V')$ such that $\mathcal{G} \setminus V_\tau$ is an $(s/2, 1)$-boundary expander as follows.

If $\mathcal{G} \setminus V_0$ is not an $(s/2, 1)$-boundary expander, then there exists a set U_1 of size $|U_1| \leq s/2$ such that $\left|\partial^{\mathcal{G}\setminus V_0}(U_1)\right| \leq |U_1|$. Delete U_1 and all its neighbours from $\mathcal{G} \setminus V_0$. If the resulting graph is not an $(s/2, 1)$-boundary expander, we repeat this process and iteratively delete vertex sets that do not satisfy the expansion condition. Formally, for $i \geq 1$ fix U_i to be any set of size $|U_i| \leq s/2$ such that

$$\left|\partial^{\mathcal{G}\setminus V_{i-1}}(U_i)\right| \leq |U_i| \ , \tag{A.7}$$

where we set

$$V_i := V_0 \cup \bigcup_{j=1}^{i} \mathcal{N}^{\mathcal{G}}(U_j) \tag{A.8}$$

(and where we note that, formally speaking, what is deleted at the ith step is $\mathcal{N}^{\mathcal{G}}(U_i)$ together with the kernel $\mathrm{Ker}(\mathcal{N}^{\mathcal{G}}(U_i))$ of this right vertex set). Since all sets U_i constructed above are non-empty, this process must terminate for some $i = \tau$ and the resulting graph $\mathcal{G} \setminus V_\tau$ is then an $(s/2, 1)$-boundary expander (note that an empty graph without vertices vacuously satisfies the expansion condition). It remains to verify that the size condition $|\mathrm{Ker}(V_\tau)| \leq |V_0|$ for the kernel of the closure of V' holds. This is immediately implied by the following inductive claim.

Claim A.1. Let $V_{-1} = U_0 = \emptyset$ and suppose that $i \geq 0$. Then for U_i satisfying (A.7) and V_i defined by (A.8) the following properties hold:

1. For all U' such that $\mathrm{Ker}(V_{i-1}) \cup U_i \subseteq U' \subseteq \mathrm{Ker}(V_i)$ we have $\left|\partial^{\mathcal{G}}(U') \setminus V_0\right| \leq |\mathrm{Ker}(V_i)|$.
2. The kernel of V_i has size $|\mathrm{Ker}(V_i)| \leq |V_0|$.

For $i = 0$, Property 1 in Claim A.1 follows because $U' \subseteq \mathrm{Ker}(V_0)$ implies that $\partial^{\mathcal{G}}(U') \subseteq V_0$. For Property 2, suppose that $|\mathrm{Ker}(V_0)| \leq s$. Then expansion implies $2|\mathrm{Ker}(V_0)| \leq |\partial^{\mathcal{G}}(\mathrm{Ker}(V_0))|$, and combining this with $\partial^{\mathcal{G}}(\mathrm{Ker}(V_0)) \subseteq V_0$ we obtain $|\mathrm{Ker}(V_0)| \leq \frac{1}{2}|V_0|$. If instead $|\mathrm{Ker}(V_0)| > s$, then we can find a subset $U' \subseteq \mathrm{Ker}(V_0)$ of size $|U'| = s$. By expansion we have $|\partial^{\mathcal{G}}(U')| \geq 2s$, which is a contradiction because as argued above we should have $\left|\partial^{\mathcal{G}}(U')\right| \leq |V_0| \leq s/2$.

For the induction step, suppose that both properties hold for $i - 1$. Let $U^* = \mathrm{Ker}(V_{i-1}) \cup U_i$ and consider any U' satisfying $U^* \subseteq U' \subseteq \mathrm{Ker}(V_i)$. We claim that every boundary element in $\partial^{\mathcal{G}}(U')$ is

either a boundary element from $\partial^{\mathcal{G}}(U^*)$ or is contained in V_0. To see this, note that since $U' \subseteq \mathrm{Ker}(V_i)$ we have $\partial^{\mathcal{G}}(U') \subseteq V_i = V_0 \cup \bigcup_{j=1}^{i} \mathcal{N}^{\mathcal{G}}(U_j)$. Furthermore, it can be observed that $\bigcup_{j=1}^{i} U_j \subseteq U^* \subseteq U'$ (this is basically due to the fact that $\mathcal{N}(\mathrm{Ker}(V')) \subseteq V'$ for any V'). Hence, if $v \in \partial^{\mathcal{G}}(U') \setminus V_0$, then it must hold that $v \in \bigcup_{j=1}^{i} \mathcal{N}^{\mathcal{G}}(U_j)$, and so the unique neighbour of v on the left is contained in $\bigcup_{j=1}^{i} U_j$ and therefore also in U^*, implying that $v \in \partial(U^*)$. This yields that

$$\partial^{\mathcal{G}}(U') \setminus V_0 \subseteq \partial^{\mathcal{G}}(U^*) \setminus V_0 \tag{A.9}$$

as claimed, and in what follows we will show

$$\left|\partial^{\mathcal{G}}(U^*) \setminus V_0\right| = \left|\partial^{\mathcal{G}}(\mathrm{Ker}(V_{i-1}) \cup U_i) \setminus V_0\right| \leq |\mathrm{Ker}(V_i)| \tag{A.10}$$

in order to prove Property 1.

Note that by construction every vertex in $V_{i-1} \setminus V_0$ has at least one neighbour in $\mathrm{Ker}(V_{i-1})$. It follows that all new boundary vertices in $\partial^{\mathcal{G}}(\mathrm{Ker}(V_{i-1}) \cup U_i) \setminus \partial^{\mathcal{G}}(\mathrm{Ker}(V_{i-1}))$ are either from V_0 or from the boundary $\partial^{\mathcal{G}\setminus V_{i-1}}(U_i)$ of U_i outside of V_{i-1}. Therefore we have

$$\partial^{\mathcal{G}}(U^*) \setminus V_0 = \partial^{\mathcal{G}}\big(\mathrm{Ker}(V_{i-1}) \cup U_i\big) \setminus V_0 \subseteq \big(\partial^{\mathcal{G}}(\mathrm{Ker}(V_{i-1})) \setminus V_0\big) \dot\cup\, \partial^{\mathcal{G}\setminus V_{i-1}}(U_i) \ . \tag{A.11}$$

Since by assumption U_i does not satisfy the expansion condition we know that

$$\left|\partial^{\mathcal{G}\setminus V_{i-1}}(U_i)\right| \leq |U_i| \tag{A.12}$$

and by the inductive hypothesis concerning Property 1 we have

$$\left|\partial^{\mathcal{G}}(\mathrm{Ker}(V_{i-1})) \setminus V_0\right| \leq |\mathrm{Ker}(V_{i-1})| \ . \tag{A.13}$$

Combining (A.9) with (A.11)–(A.13) we deduce that

$$\left|\partial^{\mathcal{G}}(U') \setminus V_0\right| \leq \left|\partial^{\mathcal{G}}(\mathrm{Ker}(V_{i-1}) \cup U_i) \setminus V_0\right| \leq |\mathrm{Ker}(V_{i-1})| + |U_i| \leq |\mathrm{Ker}(V_i)| \ , \tag{A.14}$$

where the final inequality holds since $\mathrm{Ker}(V_{i-1})$ and U_i are disjoint subsets of $\mathrm{Ker}(V_i)$. This concludes the inductive step for Property 1.

To establish Property 2, assume first that $|\mathrm{Ker}(V_i)| \leq s$. Then by expansion and Property 1 applied to $U' = \mathrm{Ker}(V_i)$ we have

$$2|\mathrm{Ker}(V_i)| \leq \left|\partial^{\mathcal{G}}(\mathrm{Ker}(V_i))\right| \leq |V_0| + |\mathrm{Ker}(V_i)| \tag{A.15}$$

and hence

$$|\mathrm{Ker}(V_i)| \leq |V_0| \tag{A.16}$$

as desired. If instead $|\mathrm{Ker}(V_i)| > s$, then by the inductive hypothesis we know that $|\mathrm{Ker}(V_{i-1})| \leq s/2$ and by construction we have $|U_i| \leq s/2$. Therefore we can find a set U' of size $|U'| = s$ satisfying the condition $\mathrm{Ker}(V_{i-1}) \cup U_i \subseteq U' \subseteq \mathrm{Ker}(V_i)$ in Property 1. From the expansion properties of $\mathcal{G}$ we conclude that $|\partial(U')| \geq 2s$, which is a contradiction because for sets U' satisfying the conditions in Property 1 we derived (A.14), which implies that $|\partial(U')| \leq |V_0| + |\mathrm{Ker}(V_{i-1})| + |U_i| \leq 3s/2$. □